



(19) **United States**

(12) **Patent Application Publication**  
**Dunne et al.**

(10) **Pub. No.: US 2023/0377480 A1**

(43) **Pub. Date: Nov. 23, 2023**

(54) **METHOD AND DEVICE FOR PRESENTING  
A GUIDED STRETCHING SESSION**

*A63B 24/00* (2006.01)

*A63B 71/06* (2006.01)

(71) Applicant: **Apple Inc.**, Cupertino, CA (US)

(52) **U.S. Cl.**

CPC ..... *G09B 19/003* (2013.01); *G06T 19/00*  
(2013.01); *A63B 24/0075* (2013.01); *A63B*  
*71/0622* (2013.01); *G09B 5/02* (2013.01)

(72) Inventors: **James J. Dunne**, San Francisco, CA  
(US); **Adeeti V. Ullal**, Los Altos, CA  
(US); **Alexander G. Bruno**, Cupertino,  
CA (US); **Daniel M. Trietsch**, San  
Jose, CA (US); **Ioana Negoita**, San  
Jose, CA (US); **Irida Mance**, San  
Francisco, CA (US); **Matthew S.**  
**DeMers**, Lexington, KY (US); **Thomas**  
**G. Salter**, Foster City, CA (US)

(57) **ABSTRACT**

In some implementations, a method includes: while present-  
ing a 3D environment, obtaining user profile and head pose  
information for a user; determining locations for visual cues  
within the 3D environment for a first portion of a guided  
stretching session based on the user profile and head pose  
information; presenting the visual cues at the determined  
locations within the 3D environment and a directional  
indicator; and in response to detecting a change to the head  
pose information: updating a location for the directional  
indicator based on the change to the head pose information;  
and in accordance with a determination that the change to  
the head pose information satisfies a criterion associated  
with a first visual cue among the visual cues, providing at  
least one of audio, haptic, or visual feedback indicating that  
the first visual cue has been completed for the first portion  
of the guided stretching session.

(21) Appl. No.: **18/200,542**

(22) Filed: **May 22, 2023**

**Related U.S. Application Data**


(60) Provisional application No. 63/344,741, filed on May  
23, 2022.

**Publication Classification**

(51) **Int. Cl.**

*G09B 19/00* (2006.01)

*G06T 19/00* (2006.01)






Figure 1

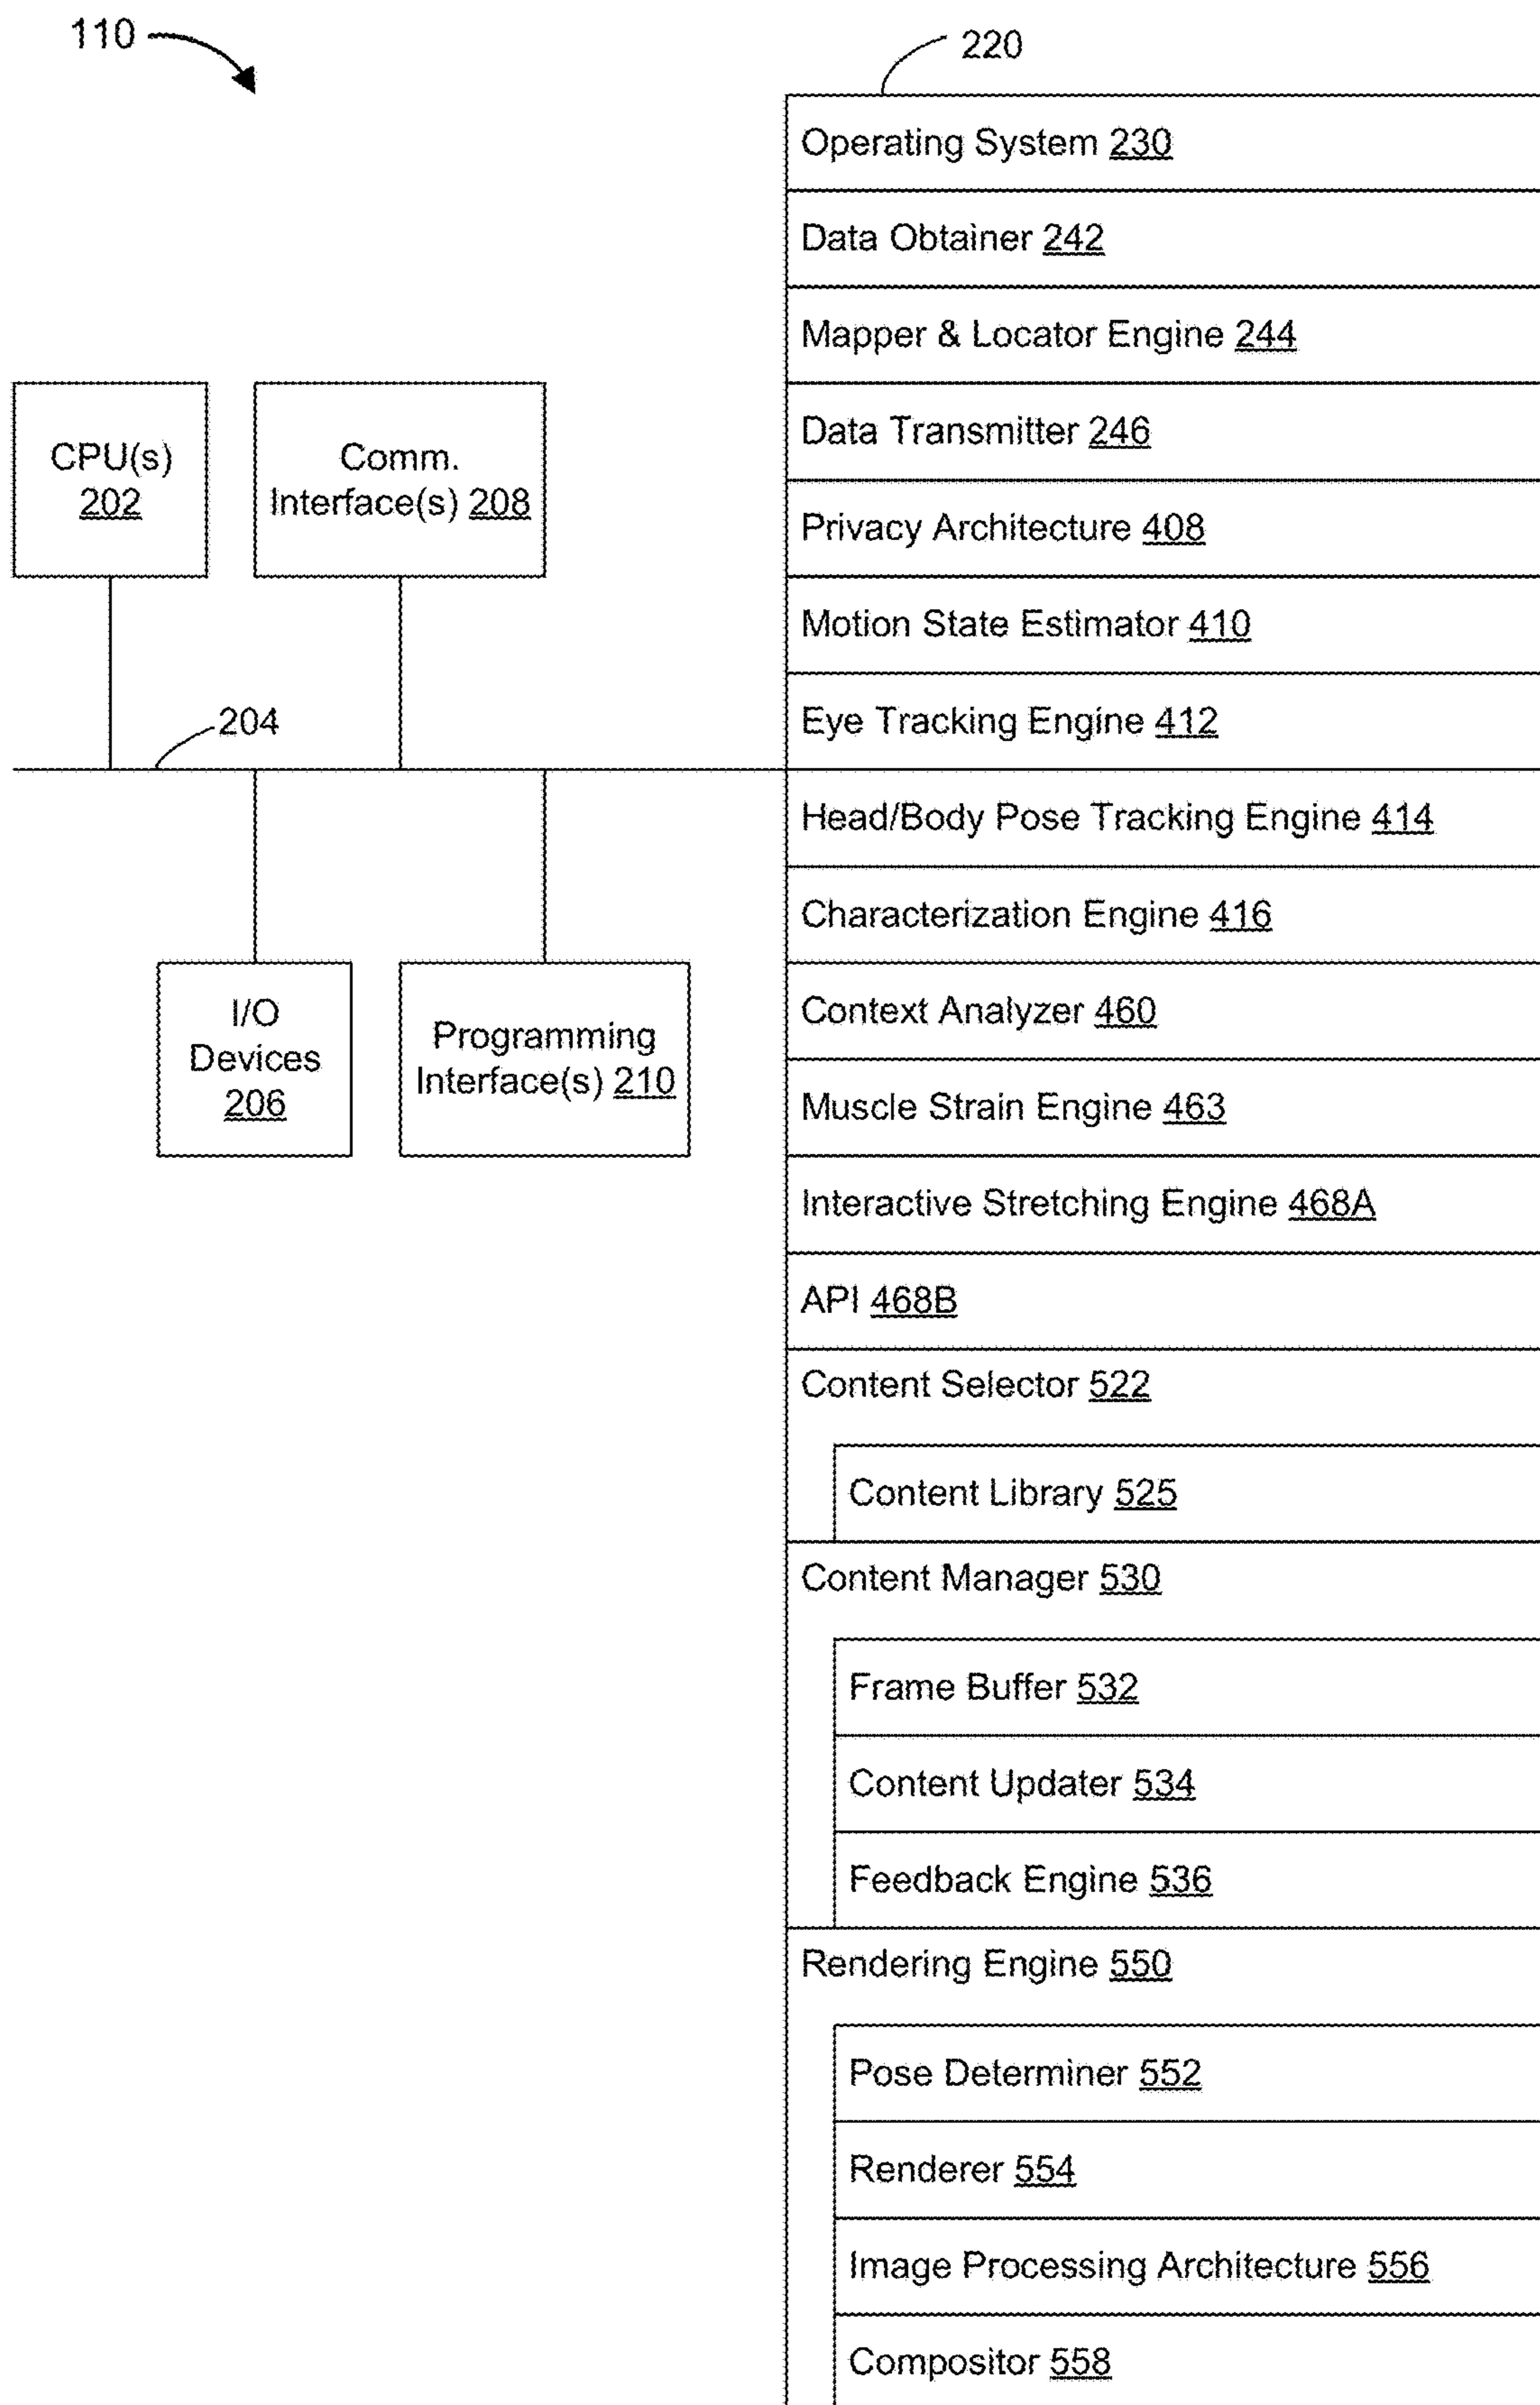


Figure 2

120 →




Figure 3



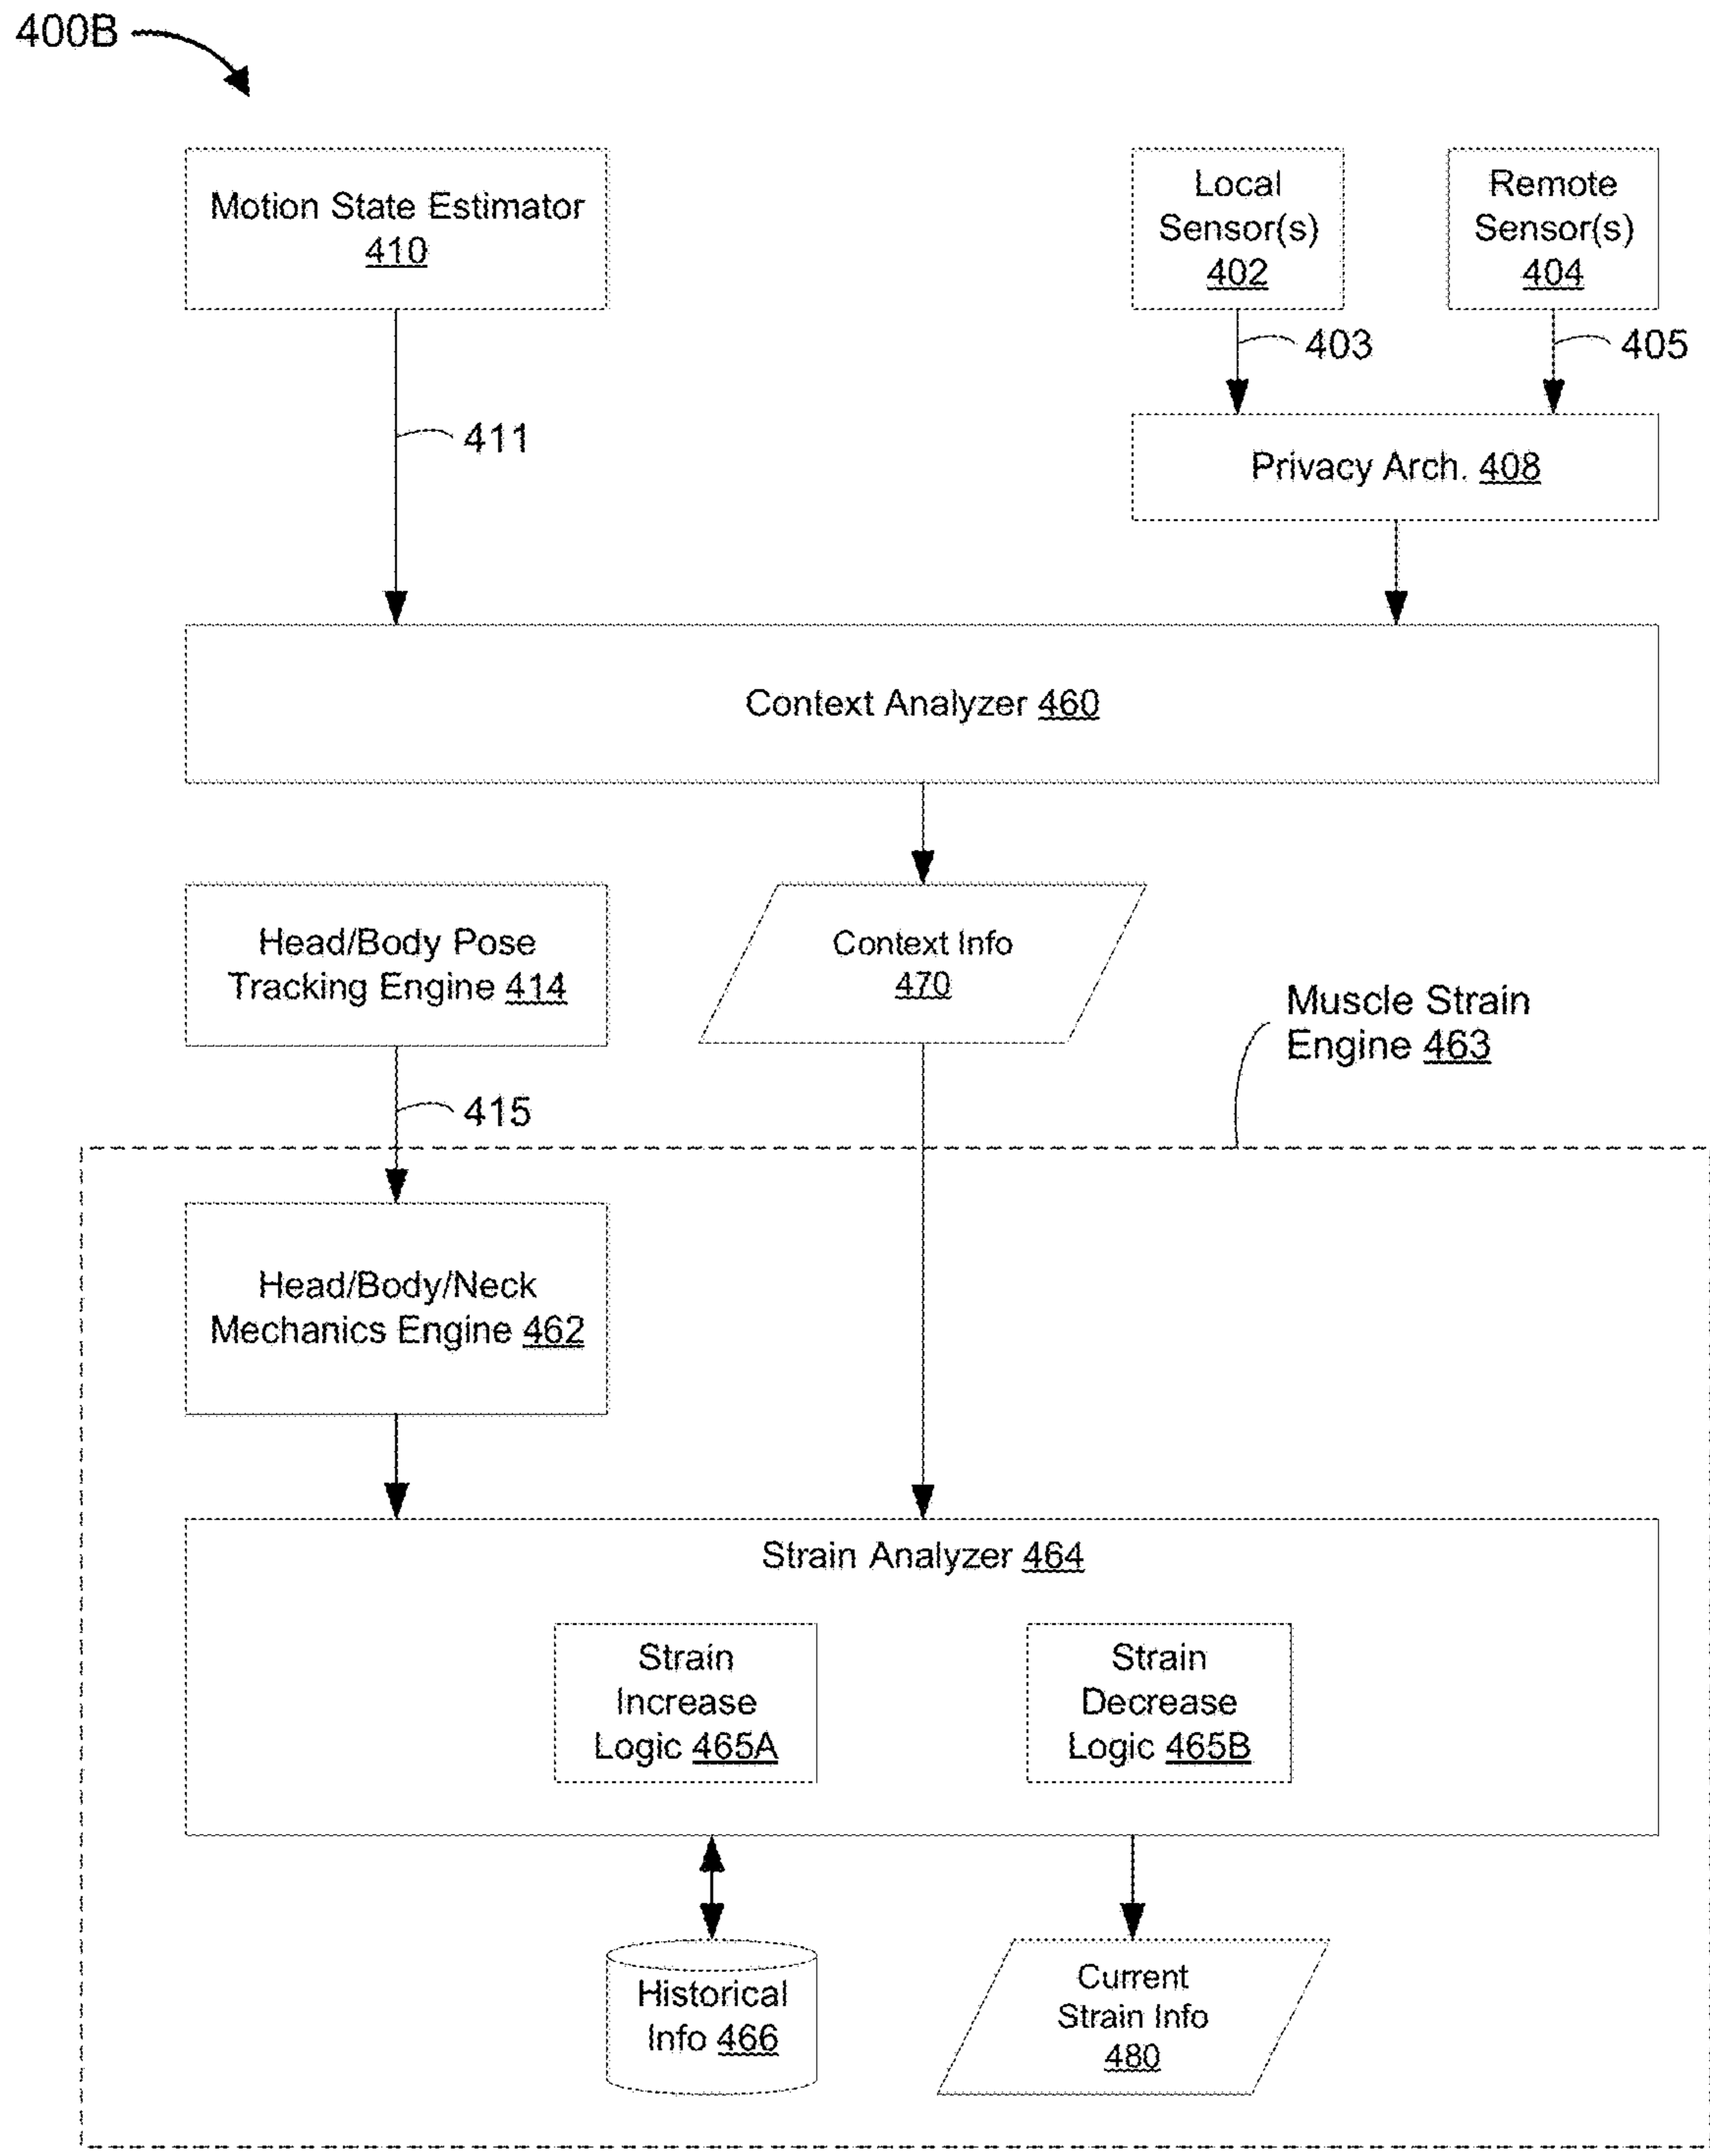


Figure 4A

**Figure 4B**





Context Info 470

|               |                    |                       |                     |           |
|---------------|--------------------|-----------------------|---------------------|-----------|
| Timestamp 471 | Env State Info 472 | Device State Info 474 | User State Info 476 | Misc. 478 |
|---------------|--------------------|-----------------------|---------------------|-----------|

Figure 4C




Figure 4D



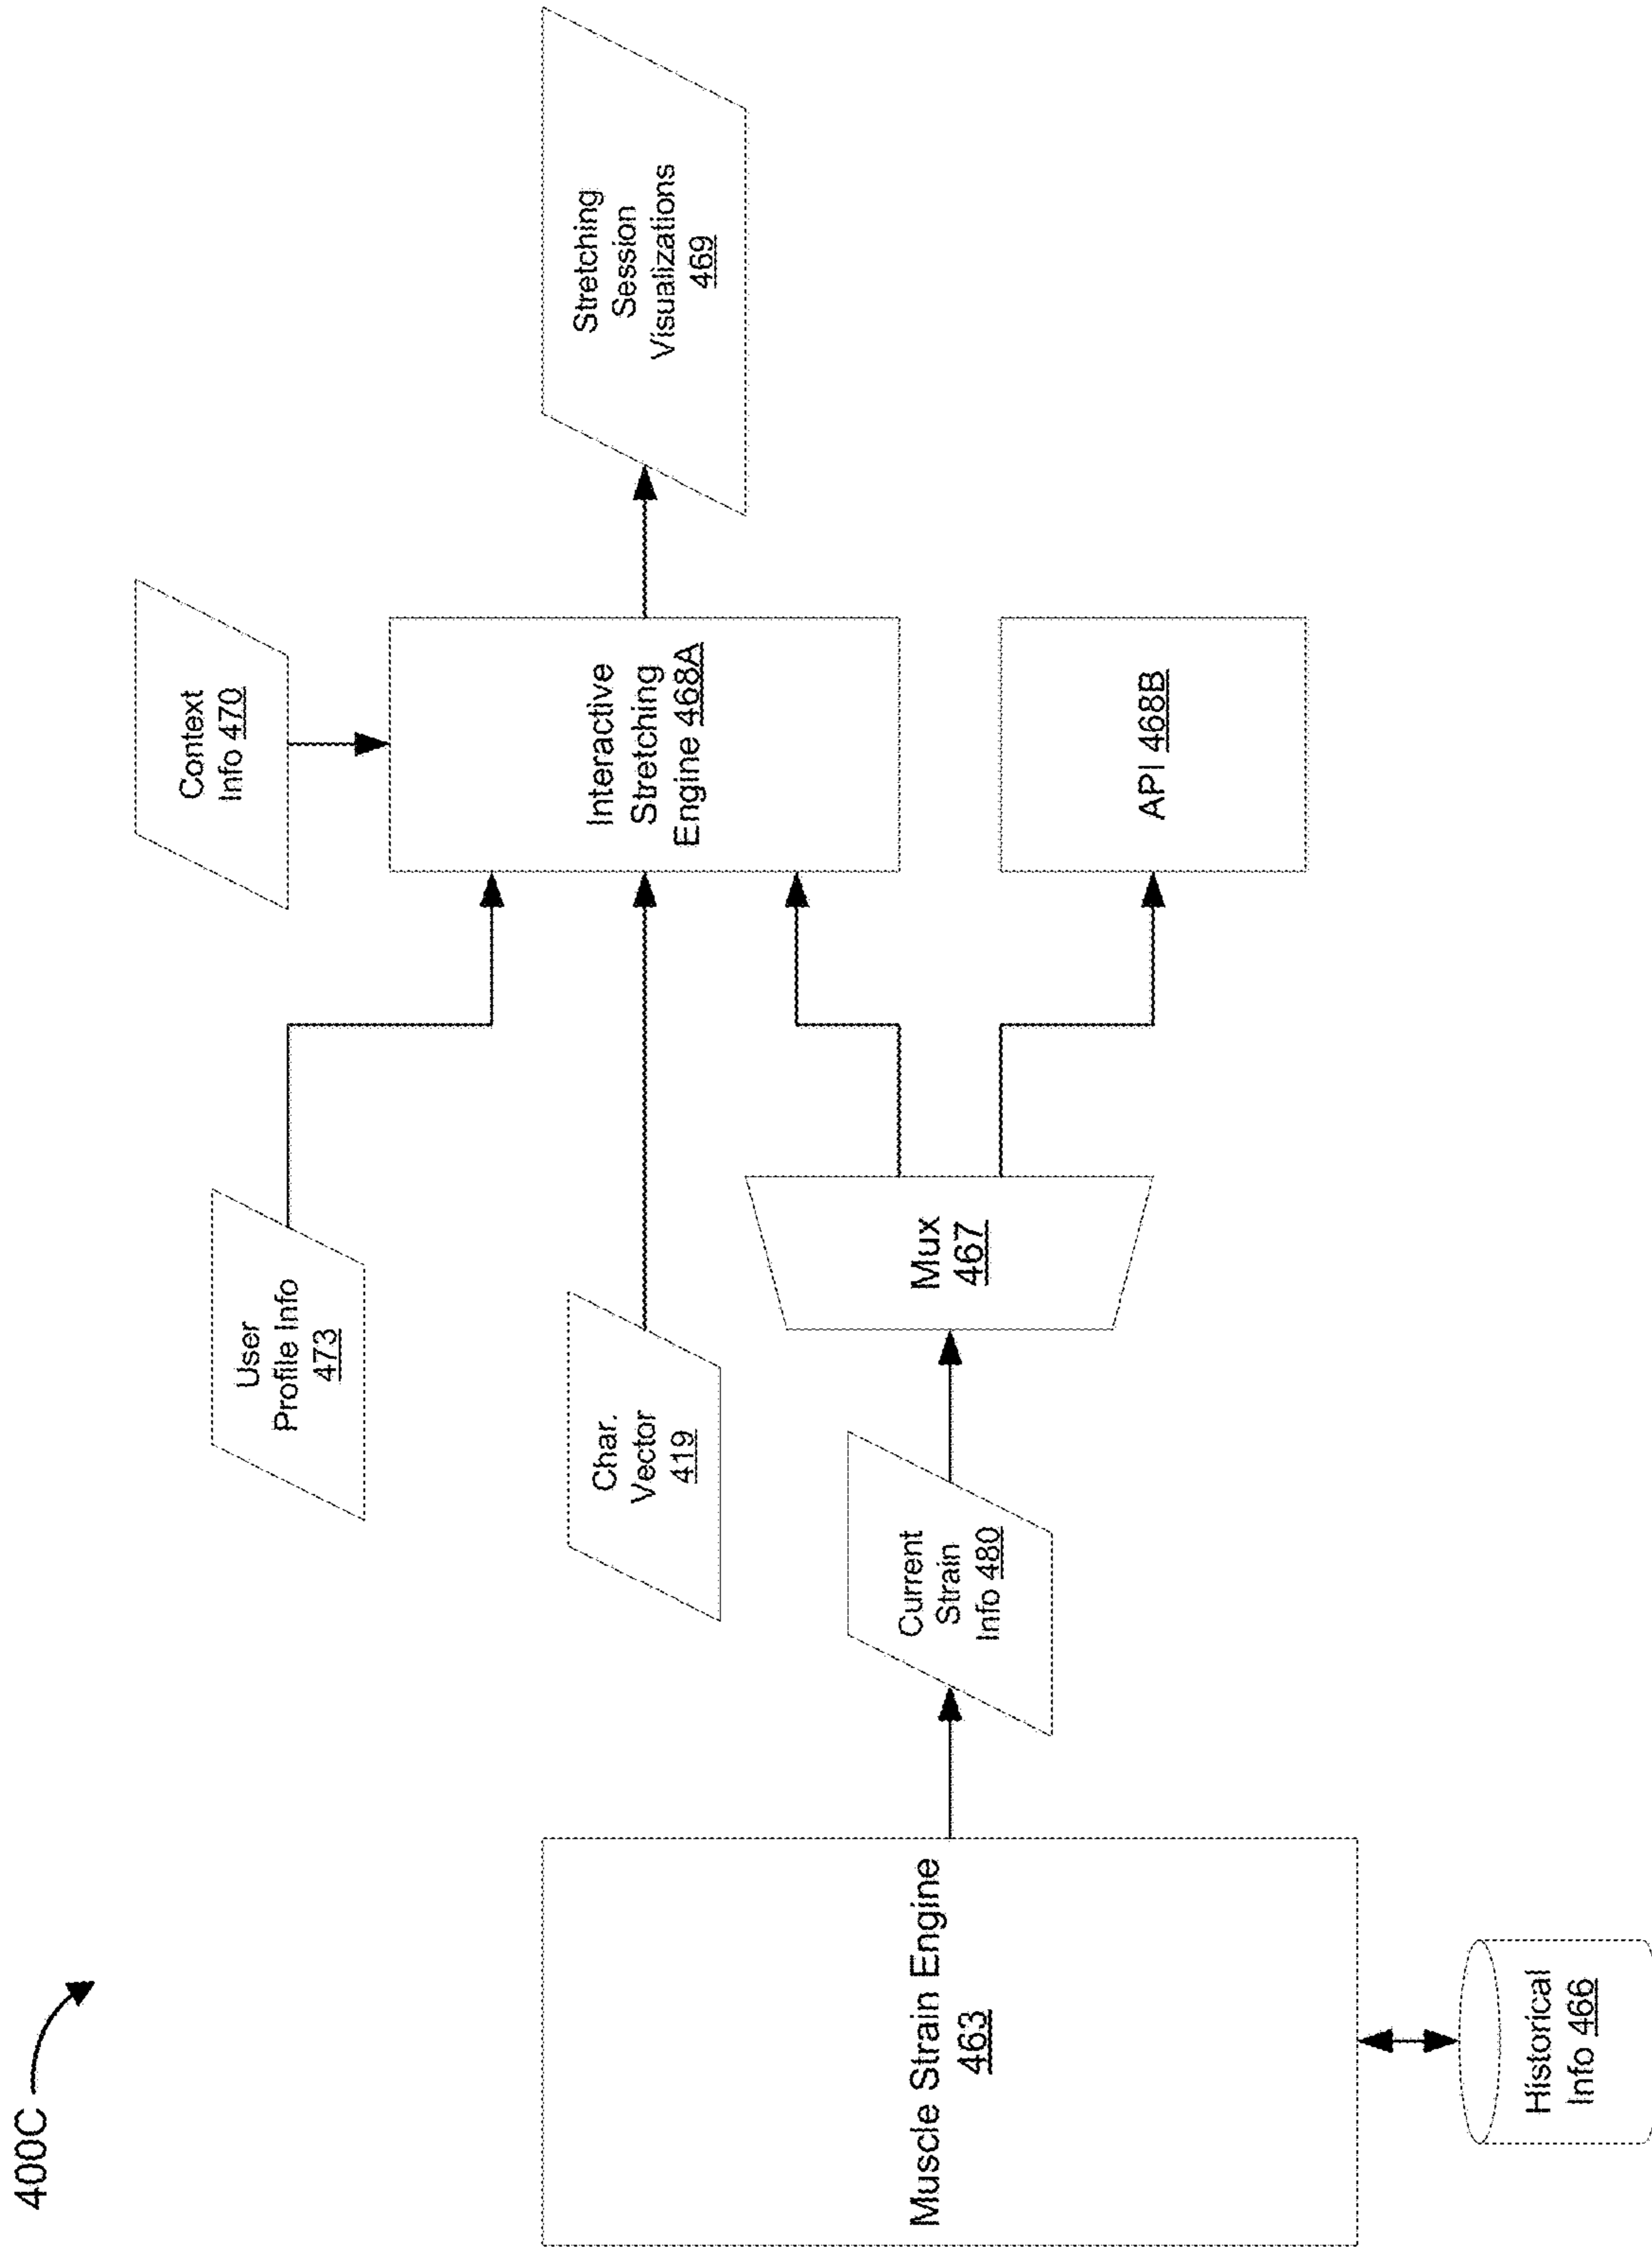


Figure 4E

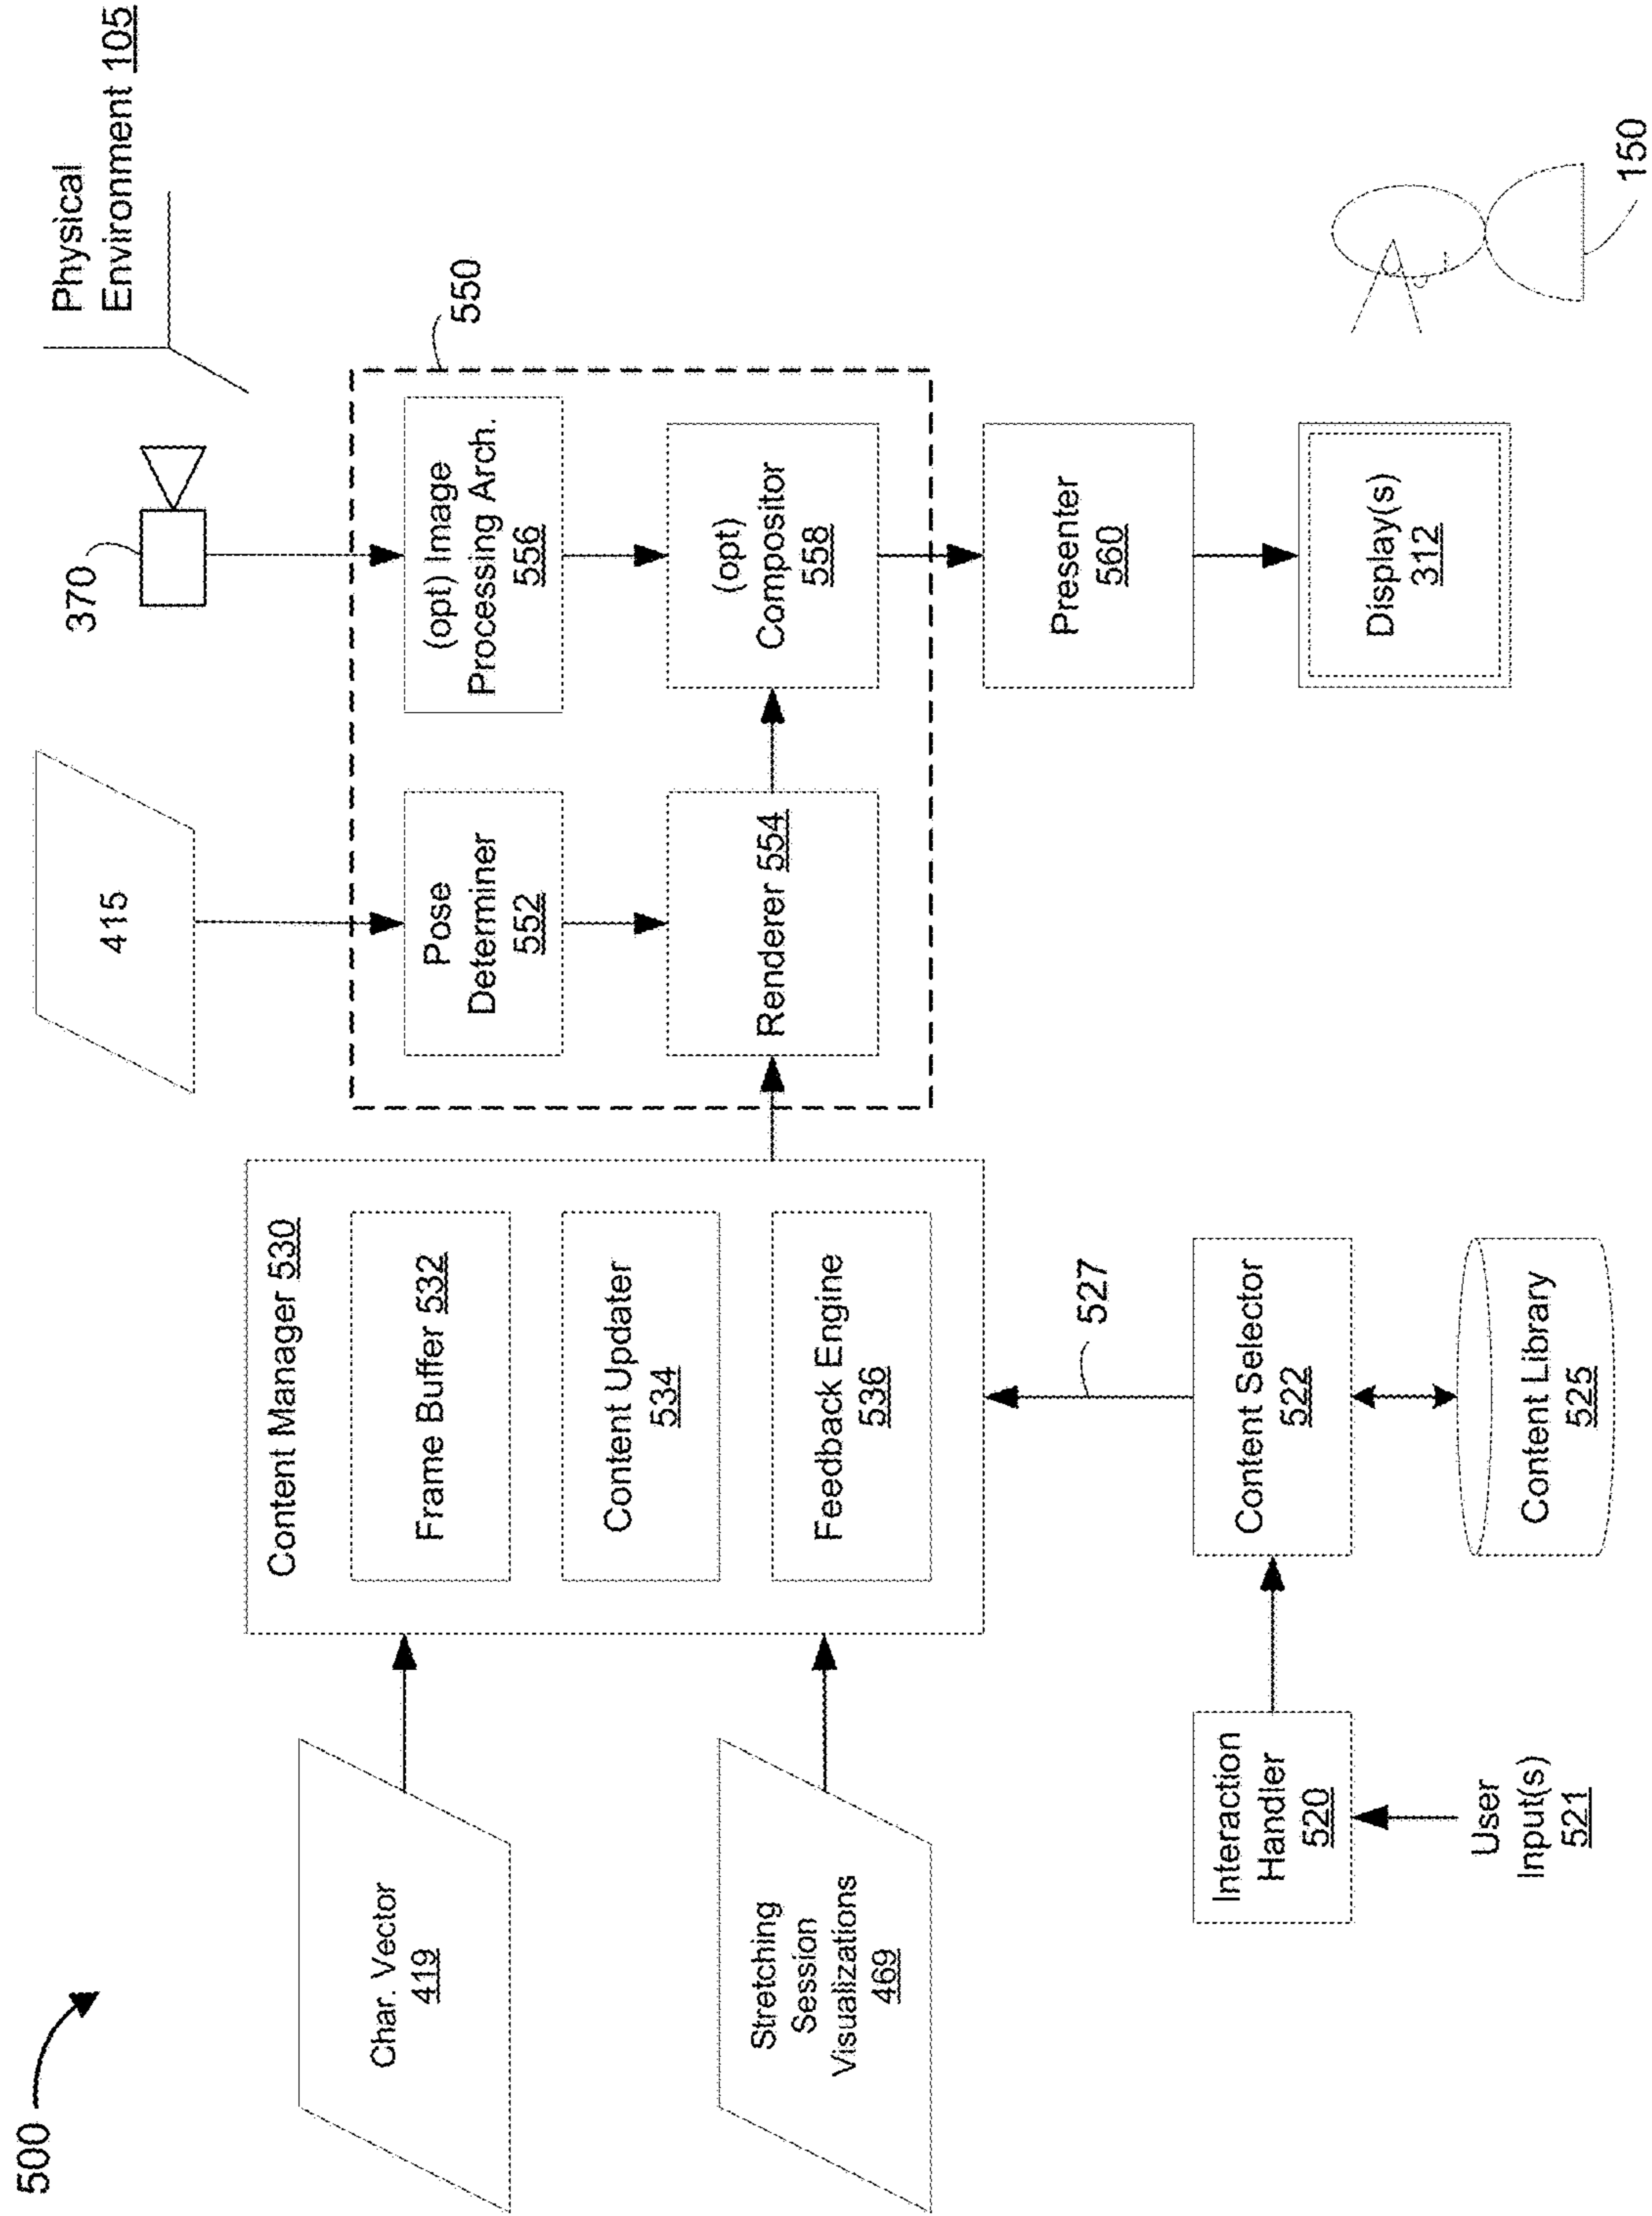


Figure 5




Figure 6A





Figure 6B






Figure 6C






Figure 6D






Figure 6E




Figure 6F

700

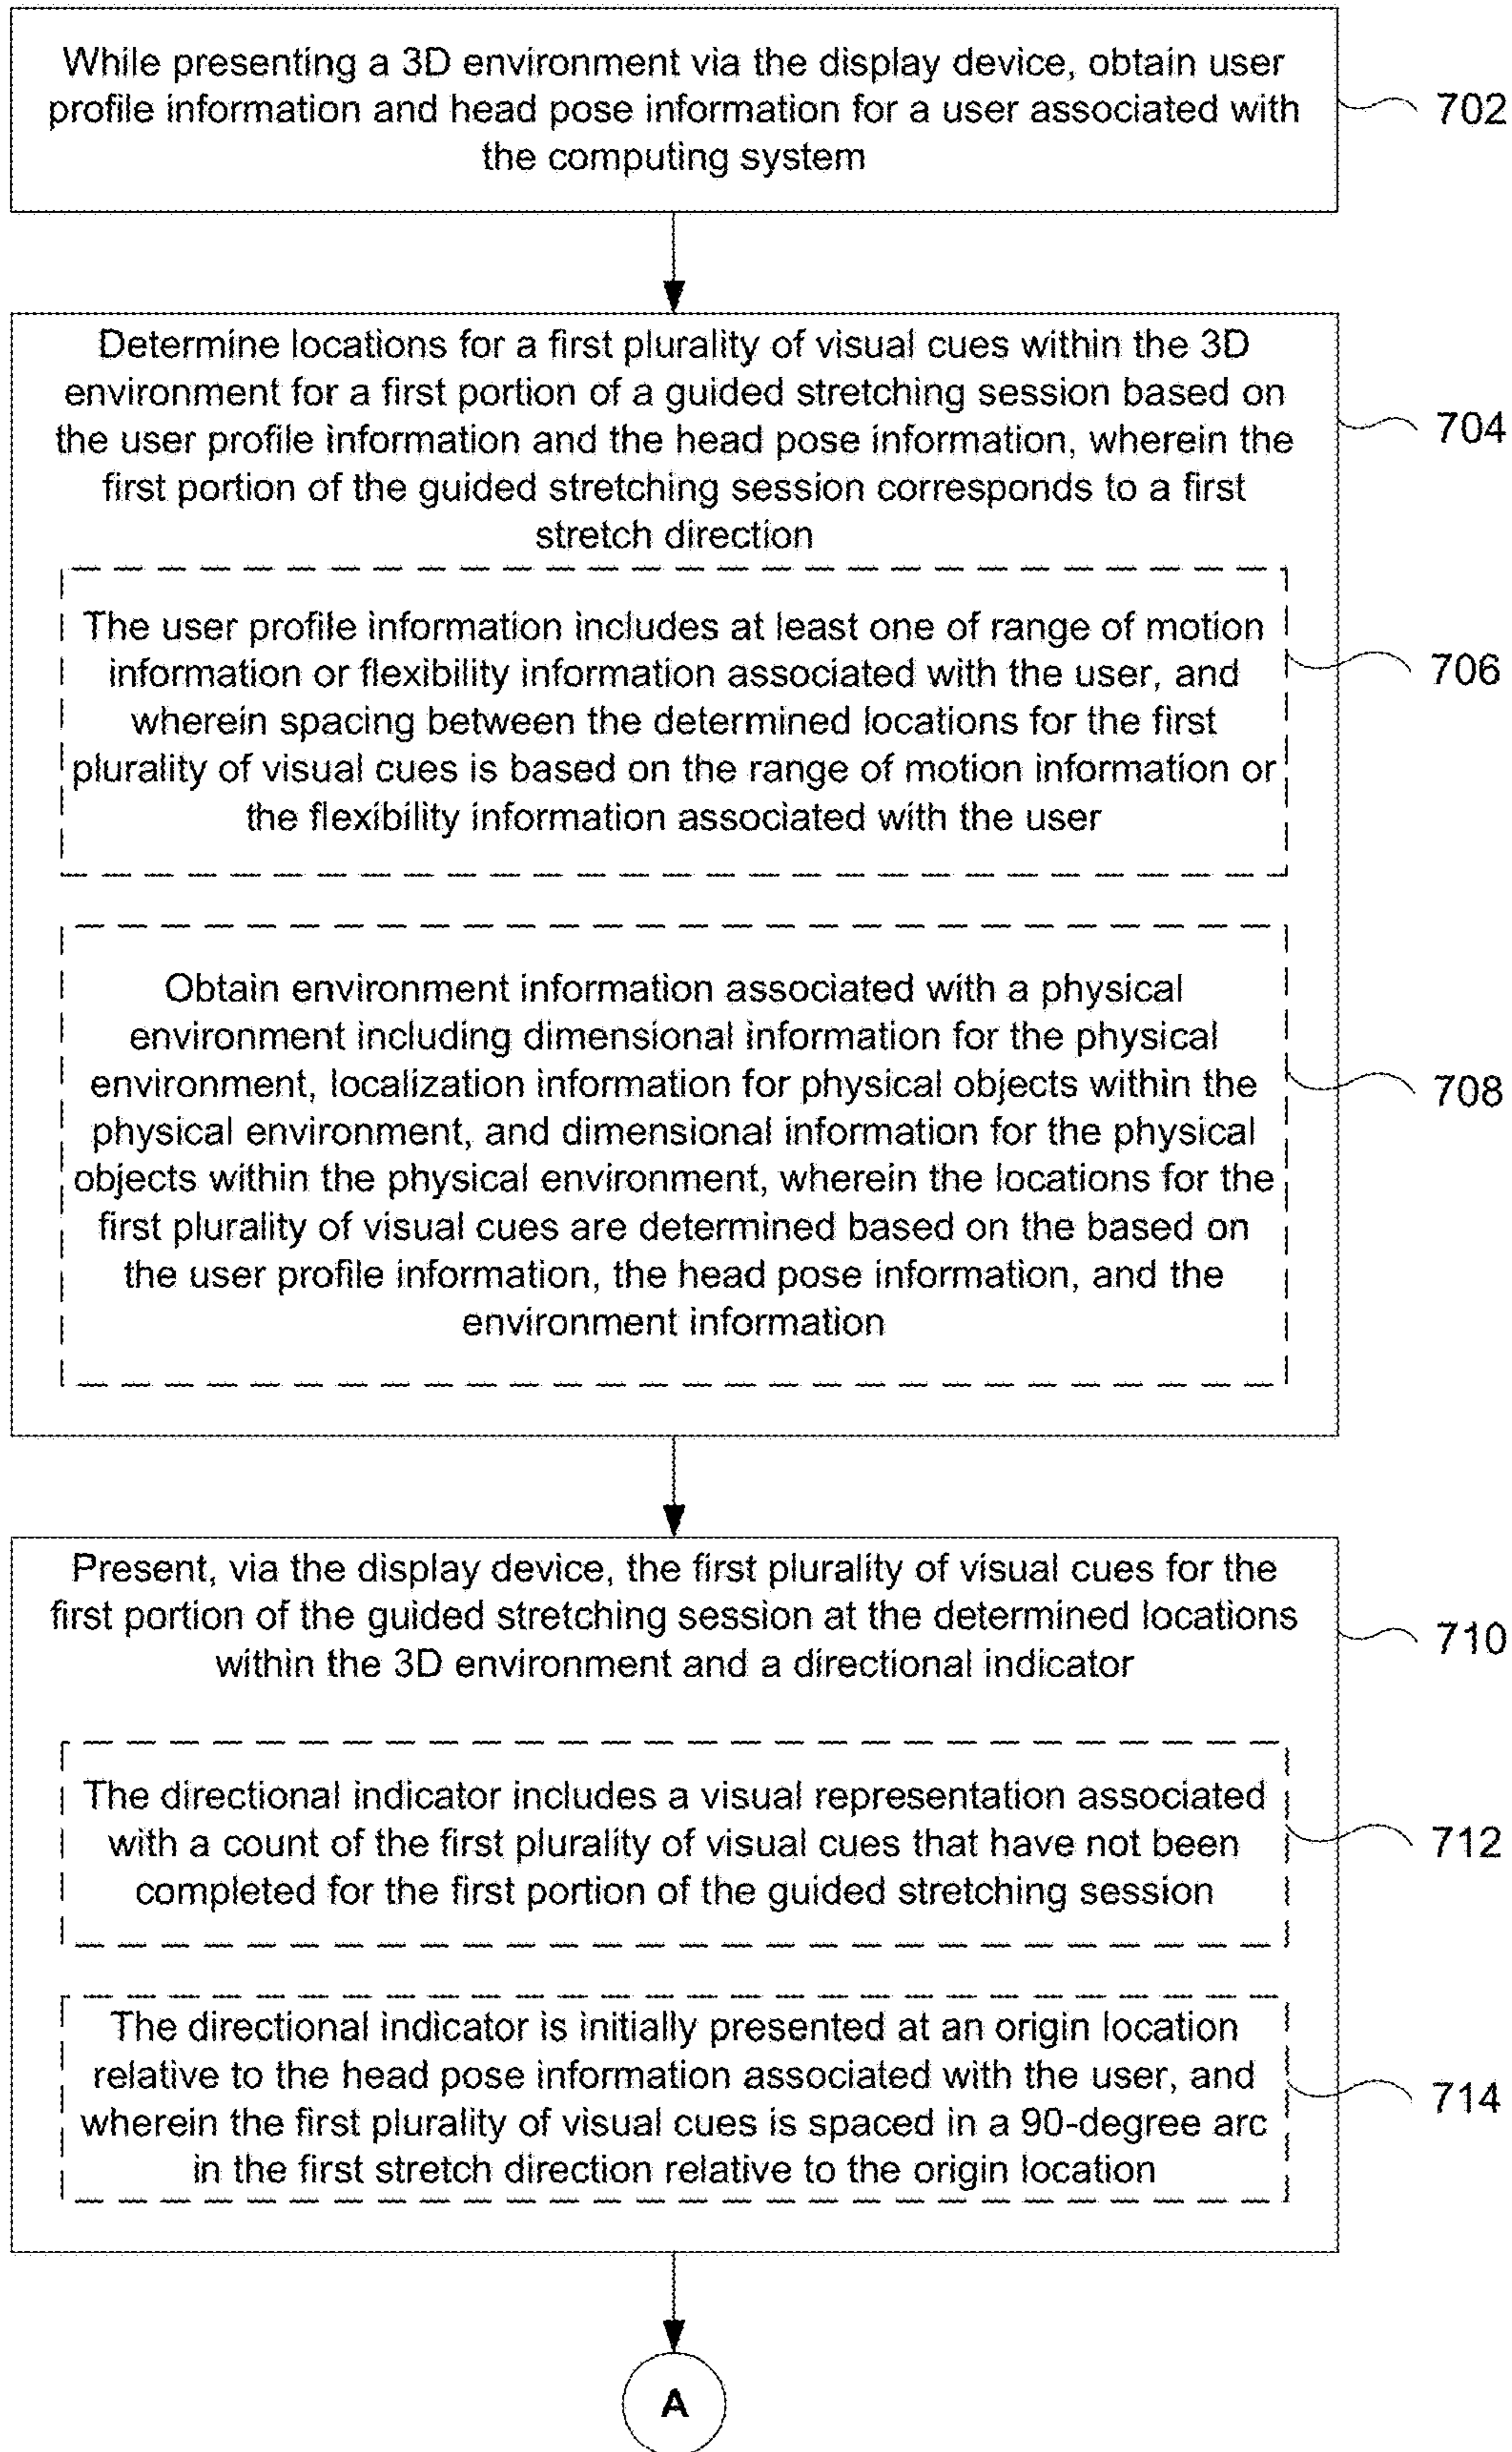


Figure 7A

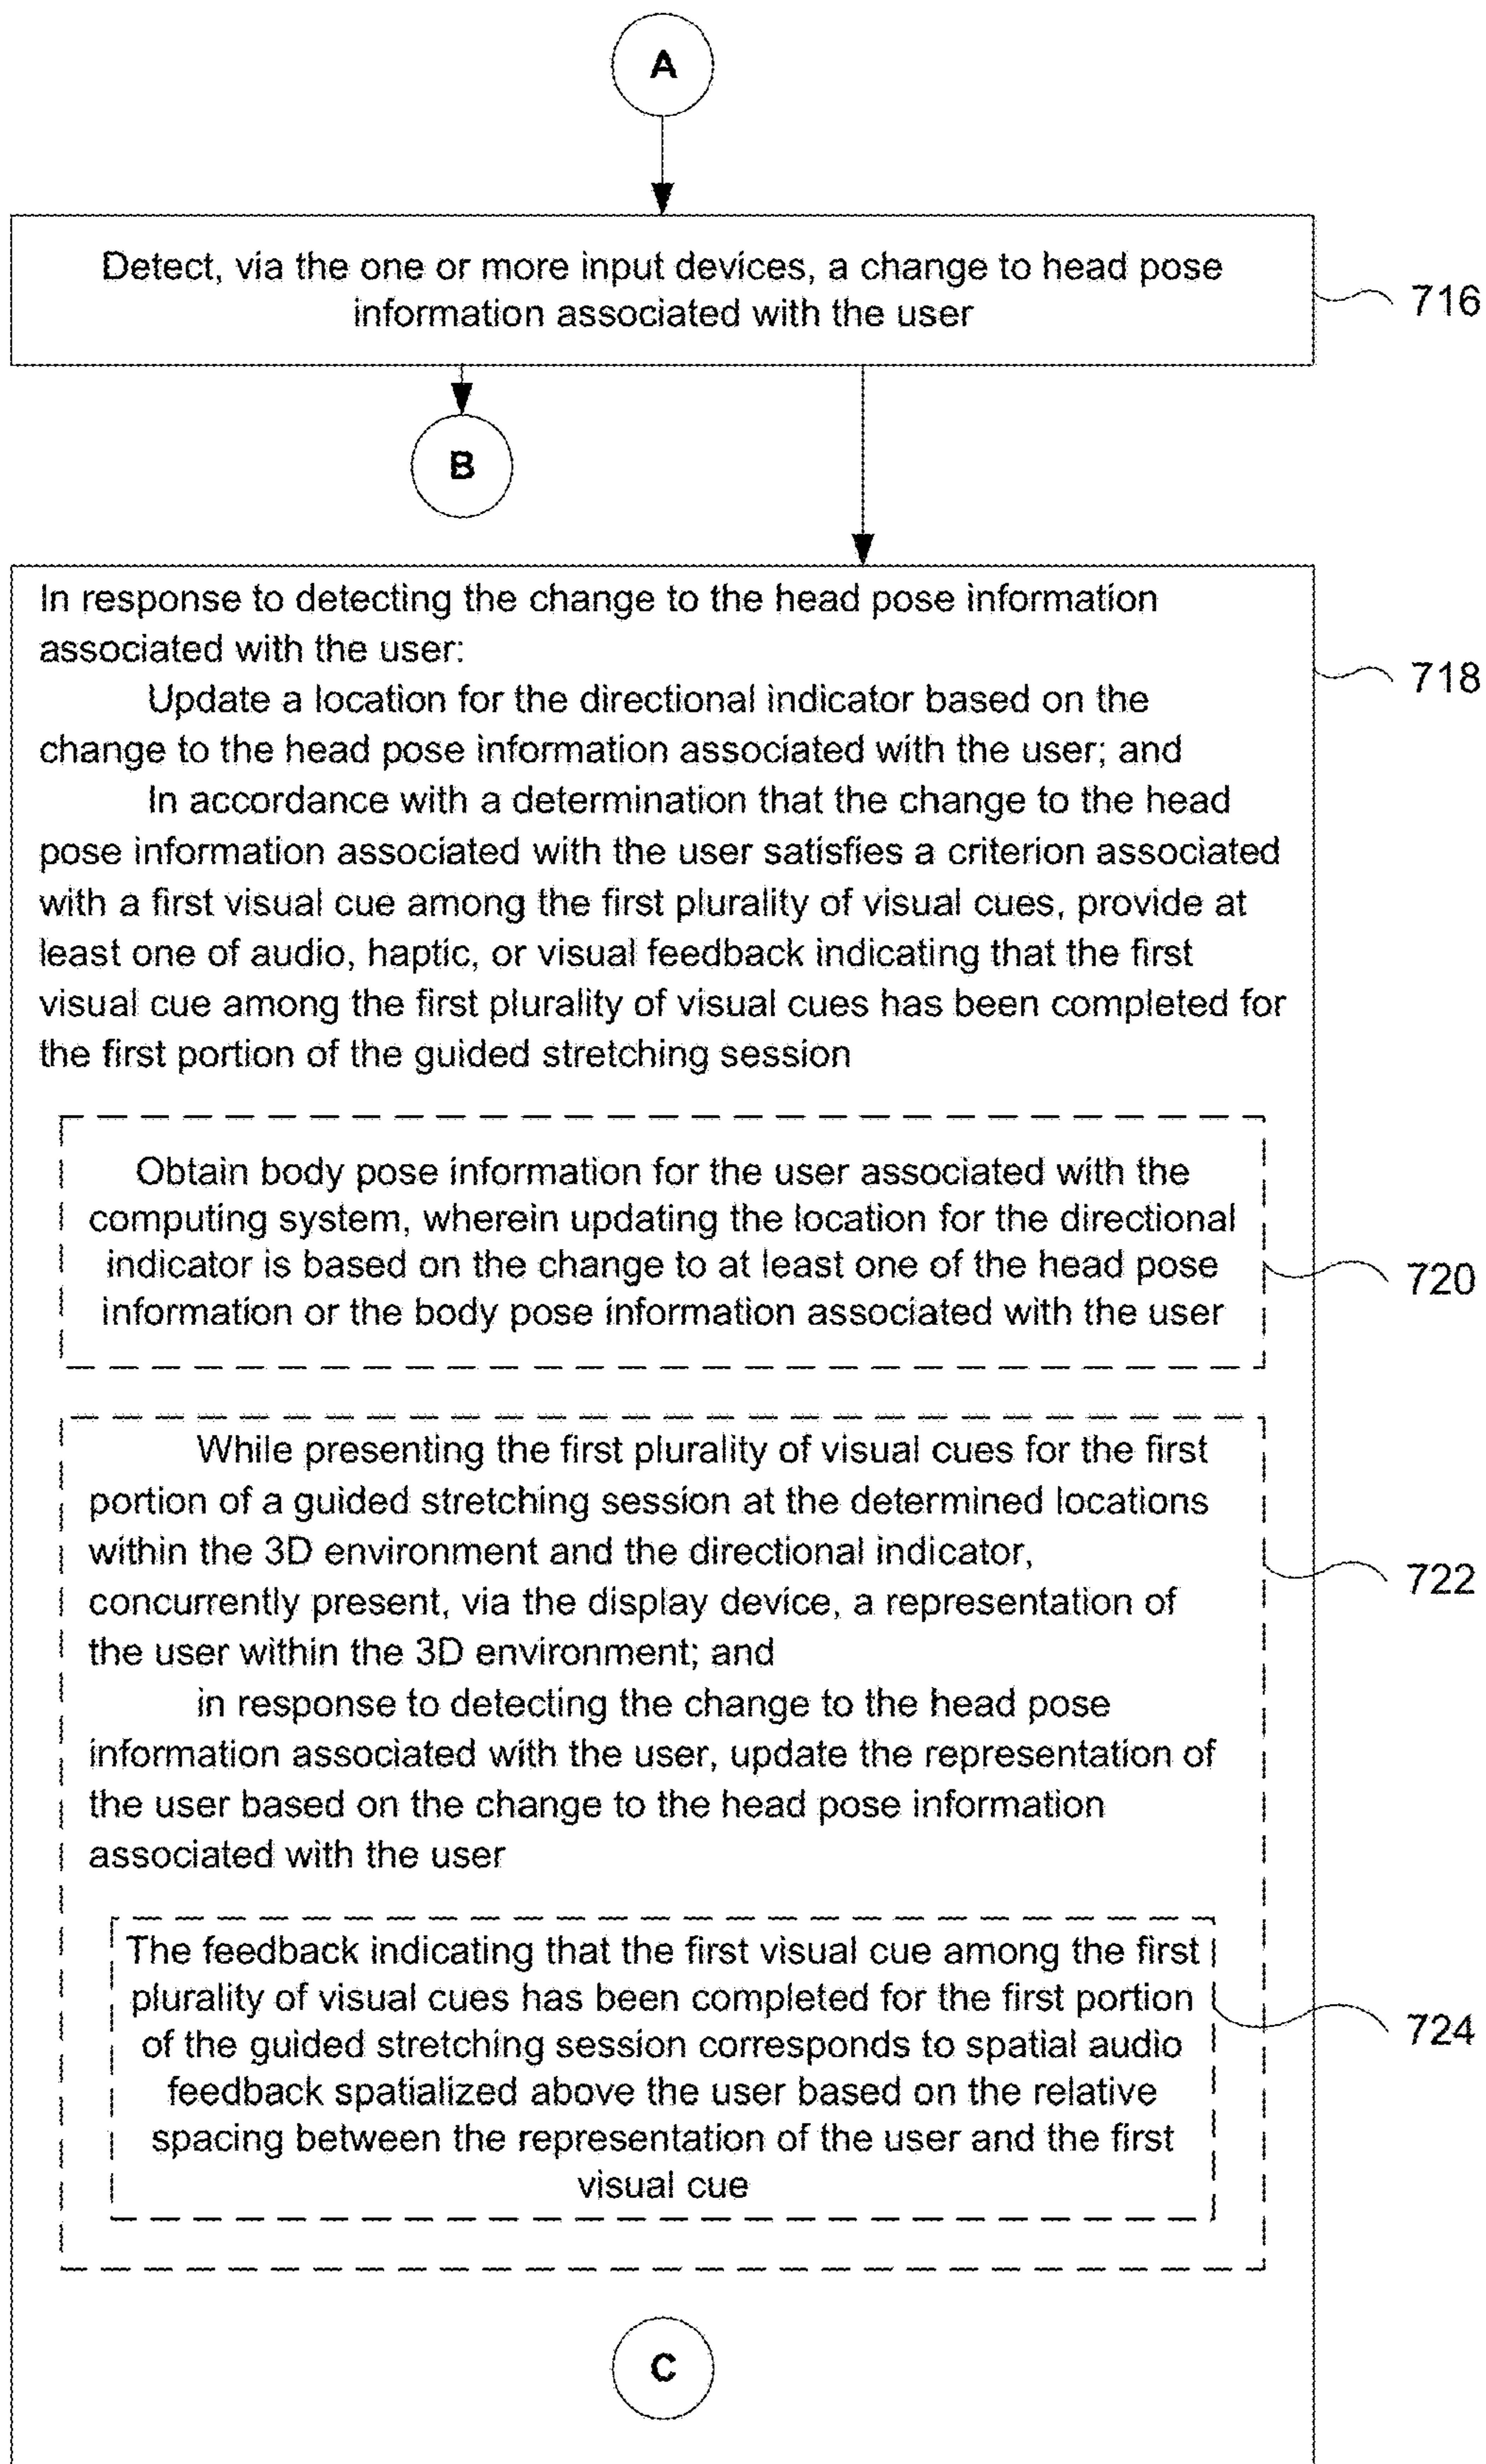


Figure 7B

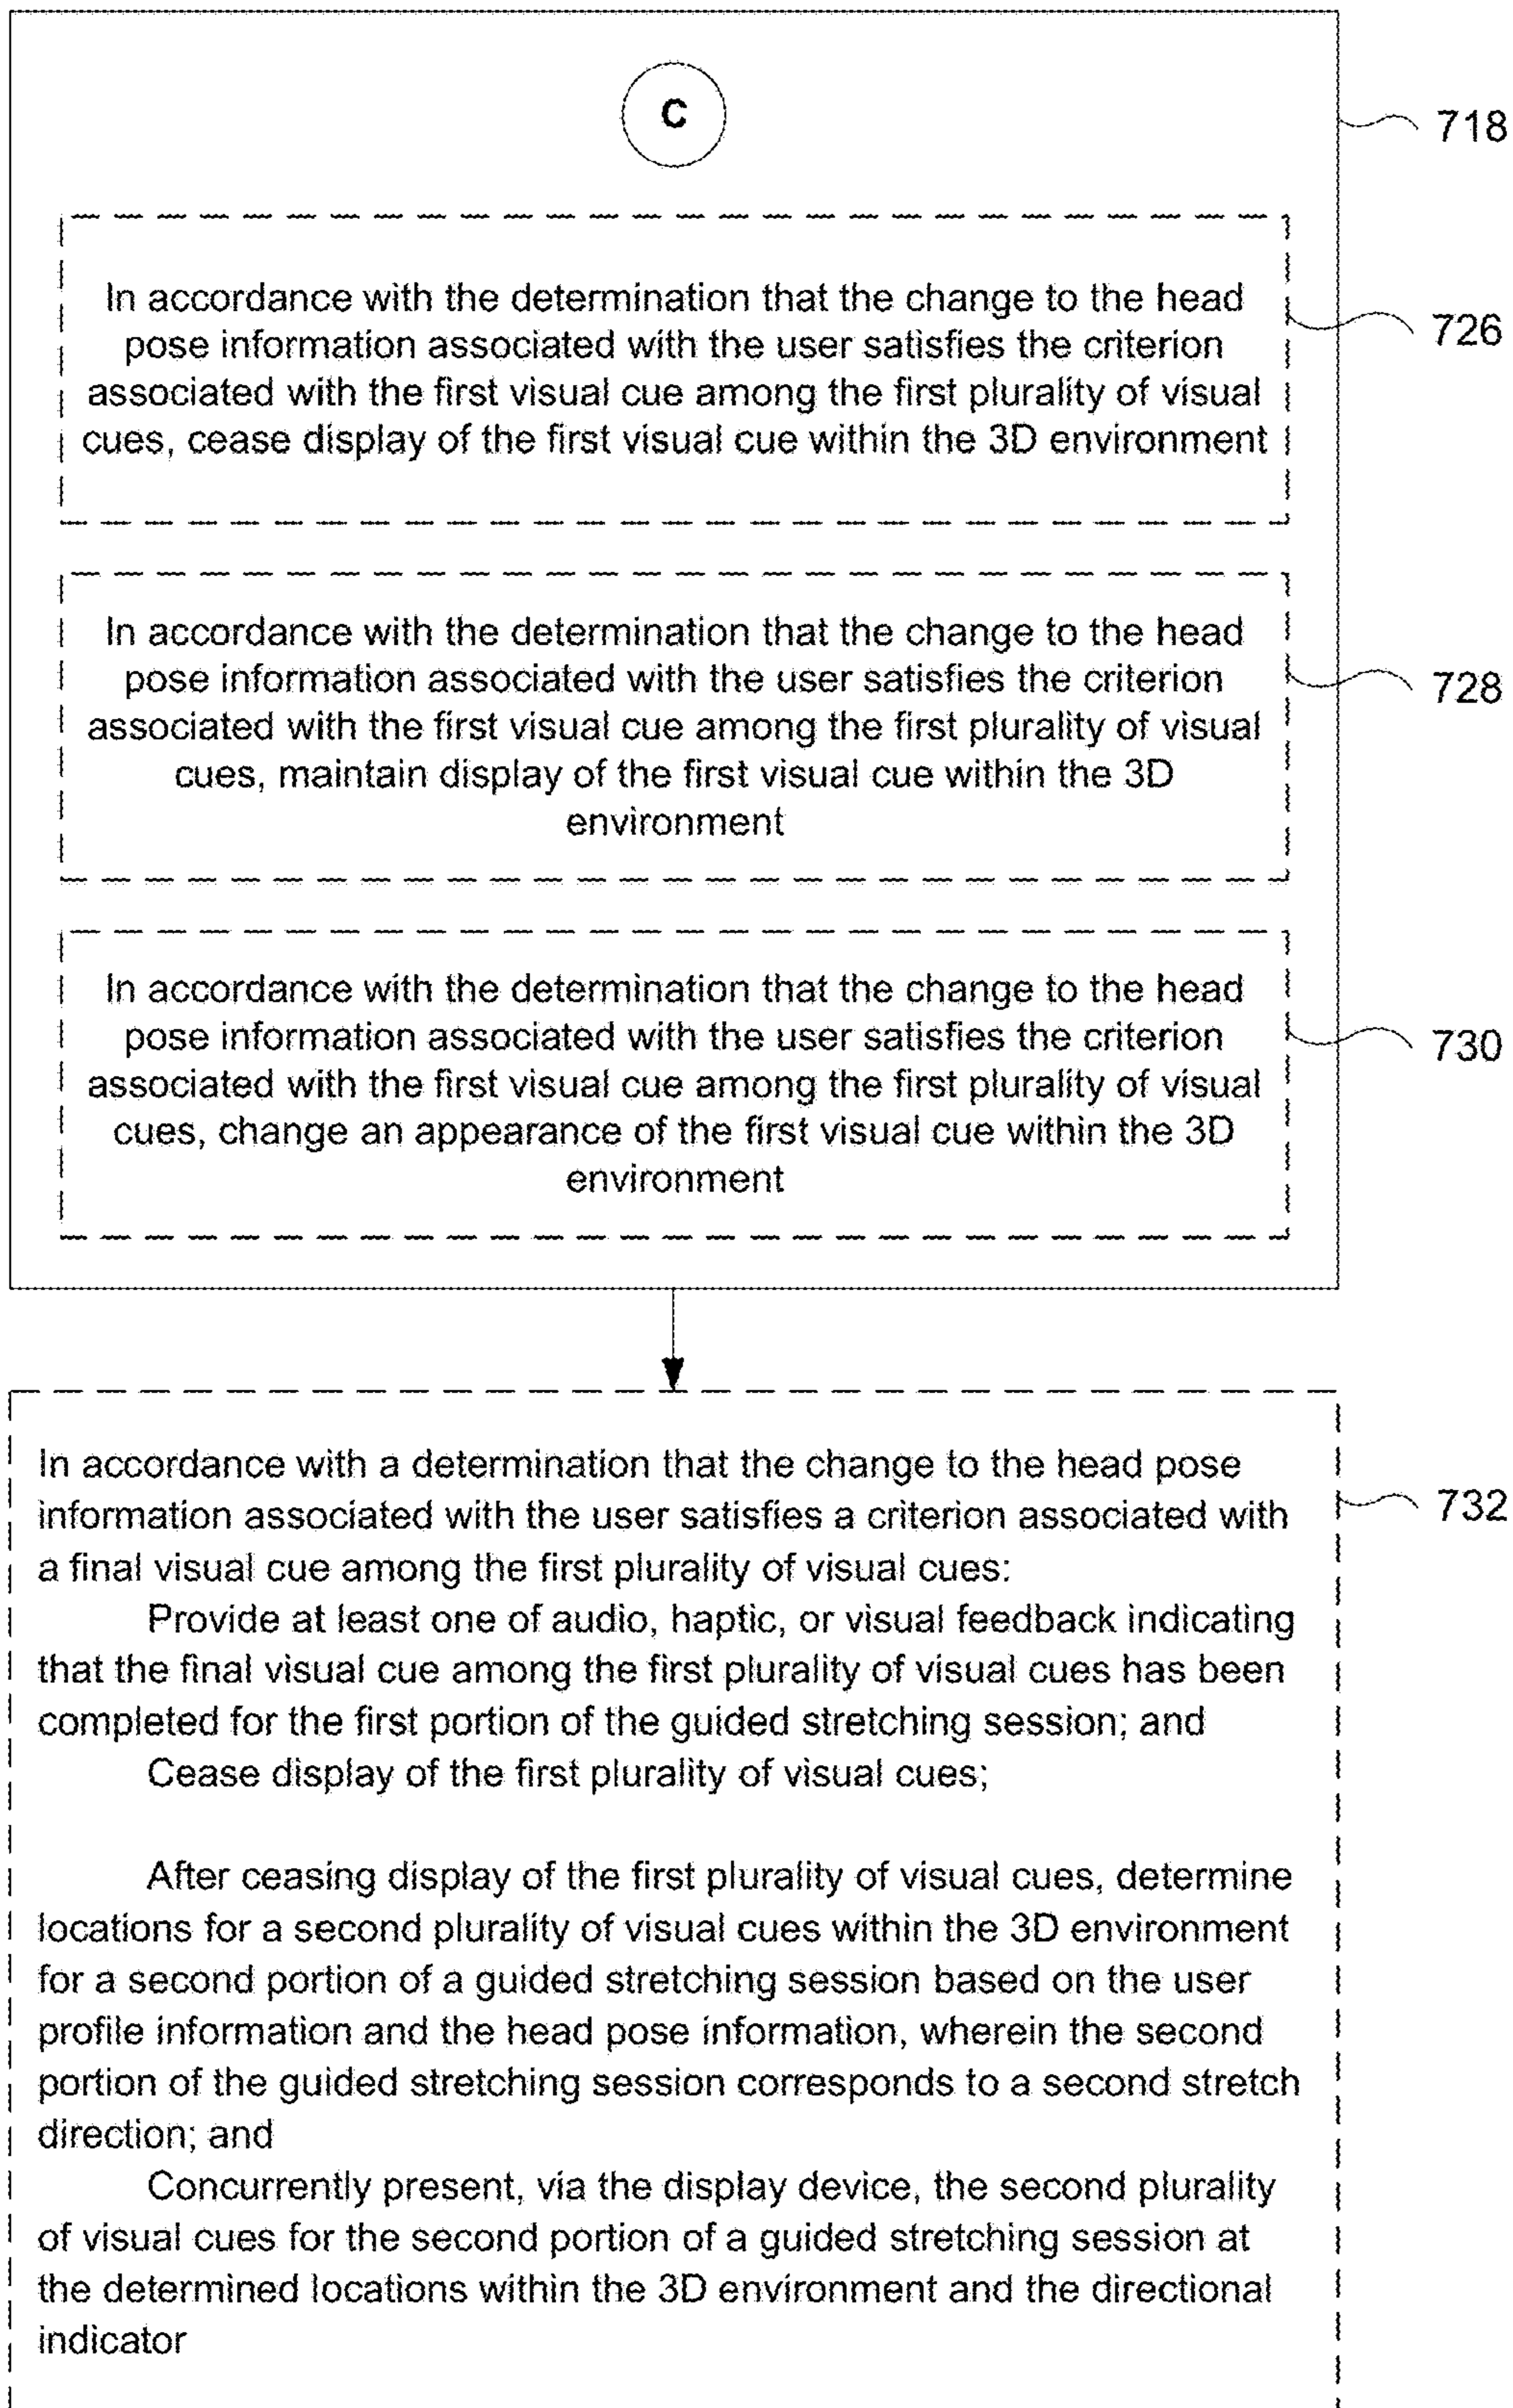



Figure 7C



**B**



In response to detecting the change to the head pose information associated with the user:  
in accordance with a determination that the change to the head pose information associated with the user does not satisfy the criterion associated with the first visual cue among the first plurality of visual cues, providing one or more hints to the user associated with the first portion of the guided stretching session

734

**Figure 7D**

## METHOD AND DEVICE FOR PRESENTING A GUIDED STRETCHING SESSION

### CROSS-REFERENCE TO RELATED APPLICATIONS

**[0001]** This application is claims priority to U.S. Provisional Patent App. No. 63/344,741, filed on May 23, 2022, which is hereby incorporated by reference in its entirety.

### TECHNICAL FIELD

**[0002]** The present disclosure generally relates to posture awareness and, in particular, to systems, devices, and methods for presenting a guided stretching session.

### BACKGROUND

**[0003]** Many persons may spend a significant number of hours at their computers or other devices during both work and non-work hours. This time spent using a computer or other device may negatively impact the posture of said person.

### BRIEF DESCRIPTION OF THE DRAWINGS

**[0004]** So that the present disclosure can be understood by those of ordinary skill in the art, a more detailed description may be had by reference to aspects of some illustrative implementations, some of which are shown in the accompanying drawings.

**[0005]** FIG. 1 is a block diagram of an example operating architecture in accordance with some implementations.

**[0006]** FIG. 2 is a block diagram of an example controller in accordance with some implementations.

**[0007]** FIG. 3 is a block diagram of an example electronic device in accordance with some implementations.

**[0008]** FIG. 4A is a block diagram of a first portion of a data processing architecture in accordance with some implementations.

**[0009]** FIG. 4B illustrates example data structures in accordance with some implementations.

**[0010]** FIG. 4C is a block diagram of a second portion of a data processing architecture in accordance with some implementations.

**[0011]** FIG. 4D illustrates example data structures in accordance with some implementations.

**[0012]** FIG. 4E is a block diagram of a third portion of a data processing architecture in accordance with some implementations.

**[0013]** FIG. 5 is a block diagram of an example content delivery architecture in accordance with some implementations.

**[0014]** FIGS. 6A-6F illustrate a plurality of three-dimensional (3D) environments associated with a guided stretching session in accordance with some implementations.

**[0015]** FIGS. 7A-7D illustrate a flowchart representation of a method of presenting a guided stretching session in accordance with some implementations.

**[0016]** In accordance with common practice the various features illustrated in the drawings may not be drawn to scale. Accordingly, the dimensions of the various features may be arbitrarily expanded or reduced for clarity. In addition, some of the drawings may not depict all of the components of a given system, method, or device. Finally, like reference numerals may be used to denote like features throughout the specification and figures.

### SUMMARY

**[0017]** Various implementations disclosed herein include devices, systems, and methods for presenting a guided stretching session. According to some implementations, the method is performed at a computing system including non-transitory memory and one or more processors, wherein the computing system is communicatively coupled to a display device and one or more input devices. The method includes: while presenting a three-dimensional (3D) environment via the display device, obtaining user profile information and head pose information for a user associated with the computing system; determining locations for a first plurality of visual cues within the 3D environment for a first portion of a guided stretching session based on the user profile information and the head pose information, wherein the first portion of the guided stretching session corresponds to a first stretch direction; presenting, via the display device, the first plurality of visual cues for the first portion of a guided stretching session at the determined locations within the 3D environment and a directional indicator; detecting, via the one or more input devices, a change to the head pose information associated with the user; and in response to detecting the change to the head pose information associated with the user: updating a location for the directional indicator based on the change to the head pose information associated with the user; and in accordance with a determination that the change to the head pose information associated with the user satisfies a criterion associated with a first visual cue among the first plurality of visual cues, providing at least one of audio, haptic, or visual feedback indicating that the first visual cue among the first plurality of visual cues has been completed for the first portion of the guided stretching session.

**[0018]** In accordance with some implementations, an electronic device includes one or more displays, one or more processors, a non-transitory memory, and one or more programs; the one or more programs are stored in the non-transitory memory and configured to be executed by the one or more processors and the one or more programs include instructions for performing or causing performance of any of the methods described herein. In accordance with some implementations, a non-transitory computer readable storage medium has stored therein instructions, which, when executed by one or more processors of a device, cause the device to perform or cause performance of any of the methods described herein. In accordance with some implementations, a device includes: one or more displays, one or more processors, a non-transitory memory, and means for performing or causing performance of any of the methods described herein.

**[0019]** In accordance with some implementations, a computing system includes one or more processors, non-transitory memory, an interface for communicating with a display device and one or more input devices, and one or more programs; the one or more programs are stored in the non-transitory memory and configured to be executed by the one or more processors and the one or more programs include instructions for performing or causing performance of the operations of any of the methods described herein. In accordance with some implementations, a non-transitory computer readable storage medium has stored therein instructions which when executed by one or more processors of a computing system with an interface for communicating with a display device and one or more input devices, cause

the computing system to perform or cause performance of the operations of any of the methods described herein. In accordance with some implementations, a computing system includes one or more processors, non-transitory memory, an interface for communicating with a display device and one or more input devices, and means for performing or causing performance of the operations of any of the methods described herein.

#### DESCRIPTION

**[0020]** Numerous details are described in order to provide a thorough understanding of the example implementations shown in the drawings. However, the drawings merely show some example aspects of the present disclosure and are therefore not to be considered limiting. Those of ordinary skill in the art will appreciate that other effective aspects and/or variants do not include all of the specific details described herein. Moreover, well-known systems, methods, components, devices, and circuits have not been described in exhaustive detail so as not to obscure more pertinent aspects of the example implementations described herein.

**[0021]** FIG. 1 is a block diagram of an example operating architecture 100 in accordance with some implementations. While pertinent features are shown, those of ordinary skill in the art will appreciate from the present disclosure that various other features have not been illustrated for the sake of brevity and so as not to obscure more pertinent aspects of the example implementations disclosed herein. To that end, as a non-limiting example, the operating architecture 100 includes an optional controller 110 and an electronic device 120 (e.g., a tablet, mobile phone, laptop, near-eye system, wearable computing device, or the like).

**[0022]** In some implementations, the controller 110 is configured to manage and coordinate an extended reality (XR) experience (sometimes also referred to herein as a “XR environment” or a “virtual environment” or a “graphical environment” or a “3D environment”) for a user 150 and optionally other users. In some implementations, the controller 110 includes a suitable combination of software, firmware, and/or hardware. The controller 110 is described in greater detail below with respect to FIG. 2. In some implementations, the controller 110 is a computing device that is local or remote relative to the physical environment 105. For example, the controller 110 is a local server located within the physical environment 105. In another example, the controller 110 is a remote server located outside of the physical environment 105 (e.g., a cloud server, central server, etc.). In some implementations, the controller 110 is communicatively coupled with the electronic device 120 via one or more wired or wireless communication channels 144 (e.g., BLUETOOTH, IEEE 802.11x, IEEE 802.16x, IEEE 802.3x, etc.). In some implementations, the functions of the controller 110 are provided by the electronic device 120. As such, in some implementations, the components of the controller 110 are integrated into the electronic device 120.

**[0023]** In some implementations, the electronic device 120 is configured to present audio and/or video (A/V) content to the user 150. In some implementations, the electronic device 120 is configured to present a user interface (UI) and/or an XR environment 128 to the user 150. In some implementations, the electronic device 120 includes a suitable combination of software, firmware, and/or hardware. The electronic device 120 is described in greater detail below with respect to FIG. 3.

**[0024]** According to some implementations, the electronic device 120 presents an XR experience to the user 150 while the user 150 is physically present within a physical environment 105 that includes a table 107 within the field-of-view (FOV) 111 of the electronic device 120. As such, in some implementations, the user 150 holds the electronic device 120 in his/her hand(s). In some implementations, while presenting the XR experience, the electronic device 120 is configured to present XR content (sometimes also referred to herein as “graphical content” or “virtual content”), including an XR cylinder 109, and to enable video pass-through of the physical environment 105 (e.g., including the table 107 (or a representations thereof) on a display 122. For example, the XR environment 128, including the XR cylinder 109, is volumetric or three-dimensional (3D).

**[0025]** In one example, the XR cylinder 109 corresponds to head/display-locked content such that the XR cylinder 109 remains displayed at the same location on the display 122 as the FOV 111 changes due to translational and/or rotational movement of the electronic device 120. As another example, the XR cylinder 109 corresponds to world/object-locked content such that the XR cylinder 109 remains displayed at its origin location as the FOV 111 changes due to translational and/or rotational movement of the electronic device 120. As such, in this example, if the FOV 111 does not include the origin location, the displayed XR environment 128 will not include the XR cylinder 109. As another example, the XR cylinder 109 corresponds to body-locked content such that it remains at a positional and rotational offset from the body of the user 150. In some examples, the electronic device 120 corresponds to a near-eye system, mobile phone, tablet, laptop, wearable computing device, or the like.

**[0026]** In some implementations, the display 122 corresponds to an additive display that enables optical see-through of the physical environment 105 including the table 107. For example, the display 122 corresponds to a transparent lens, and the electronic device 120 corresponds to a pair of glasses worn by the user 150. As such, in some implementations, the electronic device 120 presents a user interface by projecting the XR content (e.g., the XR cylinder 109) onto the additive display, which is, in turn, overlaid on the physical environment 105 from the perspective of the user 150. In some implementations, the electronic device 120 presents the user interface by displaying the XR content (e.g., the XR cylinder 109) on the additive display, which is, in turn, overlaid on the physical environment 105 from the perspective of the user 150.

**[0027]** In some implementations, the user 150 wears the electronic device 120 such as a near-eye system. As such, the electronic device 120 includes one or more displays provided to display the XR content (e.g., a single display or one for each eye). For example, the electronic device 120 encloses the FOV of the user 150. In such implementations, the electronic device 120 presents the XR environment 128 by displaying data corresponding to the XR environment 128 on the one or more displays or by projecting data corresponding to the XR environment 128 onto the retinas of the user 150.

**[0028]** In some implementations, the electronic device 120 includes an integrated display (e.g., a built-in display) that displays the XR environment 128. In some implementations, the electronic device 120 includes a head-mountable enclosure. In various implementations, the head-mountable enclosure

sure includes an attachment region to which another device with a display can be attached. For example, in some implementations, the electronic device **120** can be attached to the head-mountable enclosure. In various implementations, the head-mountable enclosure is shaped to form a receptacle for receiving another device that includes a display (e.g., the electronic device **120**). For example, in some implementations, the electronic device **120** slides/snaps into or otherwise attaches to the head-mountable enclosure. In some implementations, the display of the device attached to the head-mountable enclosure presents (e.g., displays) the XR environment **128**. In some implementations, the electronic device **120** is replaced with an XR chamber, enclosure, or room configured to present XR content in which the user **150** does not wear the electronic device **120**.

[0029] In some implementations, the controller **110** and/or the electronic device **120** cause an XR representation of the user **150** to move within the XR environment **128** based on movement information (e.g., body pose data, eye tracking data, hand/limb/finger/extremity tracking data, etc.) from the electronic device **120** and/or optional remote input devices within the physical environment **105**. In some implementations, the optional remote input devices correspond to fixed or movable sensory equipment within the physical environment **105** (e.g., image sensors, depth sensors, infrared (IR) sensors, event cameras, microphones, etc.). In some implementations, each of the remote input devices is configured to collect/capture input data and provide the input data to the controller **110** and/or the electronic device **120** while the user **150** is physically within the physical environment **105**. In some implementations, the remote input devices include microphones, and the input data includes audio data associated with the user **150** (e.g., speech samples). In some implementations, the remote input devices include image sensors (e.g., cameras), and the input data includes images of the user **150**. In some implementations, the input data characterizes body poses of the user **150** at different times. In some implementations, the input data characterizes head poses of the user **150** at different times. In some implementations, the input data characterizes hand tracking information associated with the hands of the user **150** at different times. In some implementations, the input data characterizes the velocity and/or acceleration of body parts of the user **150** such as his/her hands. In some implementations, the input data indicates joint positions and/or joint orientations of the user **150**. In some implementations, the remote input devices include feedback devices such as speakers, lights, or the like.

[0030] FIG. 2 is a block diagram of an example of the controller **110** in accordance with some implementations. While certain specific features are illustrated, those skilled in the art will appreciate from the present disclosure that various other features have not been illustrated for the sake of brevity, and so as not to obscure more pertinent aspects of the implementations disclosed herein. To that end, as a non-limiting example, in some implementations, the controller **110** includes one or more processing units **202** (e.g., microprocessors, application-specific integrated-circuits (ASICs), field-programmable gate arrays (FPGAs), graphics processing units (GPUs), central processing units (CPUs), processing cores, and/or the like), one or more input/output (I/O) devices **206**, one or more communication interfaces **208** (e.g., universal serial bus (USB), IEEE 802.3x, IEEE 802.11x, IEEE 802.16x, global system for mobile commu-

nications (GSM), code division multiple access (CDMA), time division multiple access (TDMA), global positioning system (GPS), infrared (IR), BLUETOOTH, ZIGBEE, and/or the like type interface), one or more programming (e.g., I/O) interfaces **210**, a memory **220**, and one or more communication buses **204** for interconnecting these and various other components.

[0031] In some implementations, the one or more communication buses **204** include circuitry that interconnects and controls communications between system components. In some implementations, the one or more I/O devices **206** include at least one of a keyboard, a mouse, a touchpad, a touchscreen, a joystick, one or more microphones, one or more speakers, one or more image sensors, one or more displays, and/or the like.

[0032] The memory **220** includes high-speed random-access memory such as dynamic random-access memory (DRAM), static random-access memory (SRAM), double-data-rate random-access memory (DDR RAM), or other random-access solid-state memory devices. In some implementations, the memory **220** includes non-volatile memory such as one or more magnetic disk storage devices, optical disk storage devices, flash memory devices, or other non-volatile solid-state storage devices. The memory **220** optionally includes one or more storage devices remotely located from the one or more processing units **202**. The memory **220** comprises a non-transitory computer readable storage medium. In some implementations, the memory **220** or the non-transitory computer readable storage medium of the memory **220** stores the following programs, modules and data structures, or a subset thereof described below with respect to FIG. 2.

[0033] An operating system **230** includes procedures for handling various system services and for performing hardware dependent tasks.

[0034] In some implementations, a data obtainer **242** is configured to obtain data (e.g., captured image frames of the physical environment **105**, presentation data, input data, user interaction data, camera pose tracking information, eye tracking information, head/body pose tracking information, hand/limb/finger/extremity tracking information, sensor data, location data, etc.) from at least one of the I/O devices **206** of the controller **110**, the I/O devices and sensors **306** of the electronic device **120**, and the optional remote input devices. To that end, in various implementations, the data obtainer **242** includes instructions and/or logic therefor, and heuristics and metadata therefor.

[0035] In some implementations, a mapper and locator engine **244** is configured to map the physical environment **105** and to track the position/location of at least the electronic device **120** or the user **150** with respect to the physical environment **105**. To that end, in various implementations, the mapper and locator engine **244** includes instructions and/or logic therefor, and heuristics and metadata therefor.

[0036] In some implementations, a data transmitter **246** is configured to transmit data (e.g., presentation data such as rendered image frames associated with the XR environment, location data, etc.) to at least the electronic device **120** and optionally one or more other devices. To that end, in various implementations, the data transmitter **246** includes instructions and/or logic therefor, and heuristics and metadata therefor.

[0037] In some implementations, a privacy architecture **408** is configured to ingest data and filter user information

and/or identifying information within the data based on one or more privacy filters. The privacy architecture 408 is described in more detail below with reference to FIG. 4A. To that end, in various implementations, the privacy architecture 408 includes instructions and/or logic therefor, and heuristics and metadata therefor.

[0038] In some implementations, a motion state estimator 410 is configured to obtain (e.g., receive, retrieve, or determine/generate) a motion state vector 411 associated with the electronic device 120 (and the user 150) (e.g., including a current motion state associated with the electronic device 120) based on input data and update the motion state vector 411 over time. For example, as shown in FIG. 4B, the motion state vector 411 includes a motion state descriptor 422 for the electronic device 120 (e.g., stationary, in-motion, walking, running, cycling, operating or riding in an automobile car, operating or riding in a boat, operating or riding in a bus, operating or riding in a train, operating or riding in an aircraft, or the like), translational movement values 424 associated with the electronic device 120 (e.g., a heading, a velocity value, an acceleration value, etc.), angular movement values 426 associated with the electronic device 120 (e.g., an angular velocity value, an angular acceleration value, and/or the like for each of the pitch, roll, and yaw dimensions), and/or the like. The motion state estimator 410 is described in more detail below with reference to FIG. 4A. To that end, in various implementations, the motion state estimator 410 includes instructions and/or logic therefor, and heuristics and metadata therefor.

[0039] In some implementations, an eye tracking engine 412 is configured to obtain (e.g., receive, retrieve, or determine/generate) an eye tracking vector 413 as shown in FIG. 4B (e.g., with a gaze direction) based on the input data and update the eye tracking vector 413 over time. For example, the gaze direction indicates a point (e.g., associated with x, y, and z coordinates relative to the physical environment 105 or the world-at-large), a physical object, or a region of interest (ROI) in the physical environment 105 at which the user 150 is currently looking. As another example, the gaze direction indicates a point (e.g., associated with x, y, and z coordinates relative to the XR environment 128), an XR object, or a ROI in the XR environment 128 at which the user 150 is currently looking. The eye tracking engine 412 is described in more detail below with reference to FIG. 4A. To that end, in various implementations, the eye tracking engine 412 includes instructions and/or logic therefor, and heuristics and metadata therefor.

[0040] In some implementations, a head/body pose tracking engine 414 is configured to obtain (e.g., receive, retrieve, or determine/generate) a pose characterization vector 415 based on the input data and update the pose characterization vector 415 over time. For example, as shown in FIG. 4B, the pose characterization vector 415 includes a head pose descriptor 442A (e.g., upward, downward, neutral, etc.), translational values 442B for the head pose, rotational values 442C for the head pose, a body pose descriptor 444A (e.g., standing, sitting, prone, etc.), translational values 444B for body sections/extremities/limbs/joints, rotational values 444C for the body sections/extremities/limbs/joints, and/or the like. The head/body pose tracking engine 414 is described in more detail below with reference to FIG. 4A. To that end, in various implementations, the head/body pose tracking engine 414 includes instructions and/or logic therefor, and heuristics and metadata therefor. In some imple-

mentations, the motion state estimator 410, the eye tracking engine 412, and the head/body pose tracking engine 414 may be located on the electronic device 120 in addition to or in place of the controller 110.

[0041] In some implementations, a content selector 522 is configured to select XR content (sometimes also referred to herein as “graphical content” or “virtual content”) from a content library 525 based on one or more user requests and/or inputs (e.g., a voice command, a selection from a user interface (UI) menu of XR content items or virtual agents (VAs), and/or the like). The content selector 522 is described in more detail below with reference to FIG. 4A. To that end, in various implementations, the content selector 522 includes instructions and/or logic therefor, and heuristics and metadata therefor.

[0042] In some implementations, a content library 525 includes a plurality of content items such as audio/visual (A/V) content, virtual agents (VAs), and/or XR content, objects, items, scenery, etc. As one example, the XR content includes 3D reconstructions of user captured videos, movies, TV episodes, and/or other XR content. In some implementations, the content library 525 is pre-populated or manually authored by the user 150. In some implementations, the content library 525 is located local relative to the controller 110. In some implementations, the content library 525 is located remote from the controller 110 (e.g., at a remote server, a cloud server, or the like).

[0043] In some implementations, a characterization engine 416 is configured to determine/generate a characterization vector 419 based on at least one of the motion state vector 411, the eye tracking vector 413, and the pose characterization vector 415 as shown in FIG. 4A. In some implementations, the characterization engine 416 is also configured to update the pose characterization vector 419 over time. As shown in FIG. 4B, the characterization vector 419 includes motion state information 452, gaze direction information 454, head pose information 456A, body pose information 456AB, extremity tracking information 456AC, location information 458, and/or the like. The characterization engine 416 is described in more detail below with reference to FIG. 4A. To that end, in various implementations, the characterization engine 416 includes instructions and/or logic therefor, and heuristics and metadata therefor.

[0044] In some implementations, a context analyzer 460 is configured to obtain (e.g., receive, retrieve, or determine/generate) a context information vector 470 based on input data shown in FIG. 4C and update the context information vector 470 over time. As shown in FIG. 4C, the context information vector 470 includes environment state information 472, device state information 474, and user state information 476. The context analyzer 460 is described in more detail below with reference to FIG. 4C. To that end, in various implementations, the context analyzer 460 includes instructions and/or logic therefor, and heuristics and metadata therefor.

[0045] In some implementations, a muscle strain engine 463 is configured to obtain (e.g., receive, retrieve, or determine/generate) current strain information 480 based on input data shown in FIG. 4C and update the current strain information 480 over time. As shown in FIG. 4D, the current strain information 480 includes: muscle information 482A associated with a first muscle or muscle group/region; muscle information 482B associated with a second muscle or muscle group/region; muscle information 482C associ-

ated with a third muscle or muscle group/region; muscle information **482D** associated with a fourth muscle or muscle group/region; muscle information **482E** associated with a fifth muscle or muscle group/region; and current accumulated strain information **486**. To that end, in various implementations, the muscle strain engine **463** includes a head/body/neck mechanics engine **462** and a strain analyzer **464** with strain increase logic **465A** and strain decrease logic **465B**. The muscle strain engine **463** is described in more detail below with reference to FIG. **4C**. To that end, in various implementations, the muscle strain engine **463** includes instructions and/or logic therefor, and heuristics and metadata therefor.

[0046] In some implementations, an interactive stretching engine **468A** generates stretching session visualizations **469**. The interactive stretching engine **468A** is described in more detail below with reference to FIG. **4E**. To that end, in various implementations, the interactive stretching engine **468A** includes instructions and/or logic therefor, and heuristics and metadata therefor.

[0047] In some implementations, an application programming interface (API) **468B** is configured to provide access to the current strain information **480** to at least one of: the operating system of the controller **110**, the electronic device **120**, or a combination thereof; third-party programs or applications; and/or the like. As such, the current strain information **480** may be used in various downstream processes. The API **468B** is described in more detail below with reference to FIG. **4E**. To that end, in various implementations, the API **468B** includes instructions and/or logic therefor, and heuristics and metadata therefor.

[0048] In some implementations, a content manager **530** is configured to manage and update the layout, setup, structure, and/or the like for the XR environment **128** including one or more of VA(s), XR content, one or more user interface (UI) elements associated with the XR content, and/or the like. The content manager **530** is described in more detail below with reference to FIG. **5**. To that end, in various implementations, the content manager **530** includes instructions and/or logic therefor, and heuristics and metadata therefor. In some implementations, the content manager **530** includes a frame buffer **532**, a content updater **534**, and a feedback engine **536**. In some implementations, the frame buffer **532** includes XR content, a rendered image frame, and/or the like for one or more past instances and/or frames.

[0049] In some implementations, the content updater **534** is configured to modify the XR environment **128** over time based on translational or rotational movement of the electronic device **120** or physical objects within the physical environment **105**, user inputs (e.g., a change in context, hand/extremity tracking inputs, eye tracking inputs, touch inputs, voice commands, modification/manipulation inputs with the physical object, and/or the like), and/or the like. To that end, in various implementations, the content updater **534** includes instructions and/or logic therefor, and heuristics and metadata therefor.

[0050] In some implementations, the feedback engine **536** is configured to generate sensory feedback (e.g., visual feedback such as text or lighting changes, audio feedback, haptic feedback, etc.) associated with the XR environment **128**. To that end, in various implementations, the feedback engine **536** includes instructions and/or logic therefor, and heuristics and metadata therefor.

[0051] In some implementations, a rendering engine **550** is configured to render an XR environment **128** (sometimes also referred to herein as a “graphical environment” or “virtual environment”) or image frame associated therewith as well as the VA(s), XR content, one or more UI elements associated with the XR content, and/or the like. To that end, in various implementations, the rendering engine **550** includes instructions and/or logic therefor, and heuristics and metadata therefor. In some implementations, the rendering engine **550** includes a pose determiner **552**, a renderer **554**, an optional image processing architecture **556**, and an optional compositor **558**. One of ordinary skill in the art will appreciate that the optional image processing architecture **556** and the optional compositor **558** may be present for video pass-through configurations but may be removed for fully VR or optical see-through configurations.

[0052] In some implementations, the pose determiner **552** is configured to determine a current camera pose of the electronic device **120** and/or the user **150** relative to the A/V content and/or XR content. The pose determiner **552** is described in more detail below with reference to FIG. **5**. To that end, in various implementations, the pose determiner **552** includes instructions and/or logic therefor, and heuristics and metadata therefor.

[0053] In some implementations, the renderer **554** is configured to render the A/V content and/or the XR content according to the current camera pose relative thereto. The renderer **554** is described in more detail below with reference to FIG. **5**. To that end, in various implementations, the renderer **554** includes instructions and/or logic therefor, and heuristics and metadata therefor.

[0054] In some implementations, the image processing architecture **556** is configured to obtain (e.g., receive, retrieve, or capture) an image stream including one or more images of the physical environment **105** from the current camera pose of the electronic device **120** and/or the user **150**. In some implementations, the image processing architecture **556** is also configured to perform one or more image processing operations on the image stream such as warping, color correction, gamma correction, sharpening, noise reduction, white balance, and/or the like. The image processing architecture **556** is described in more detail below with reference to FIG. **5**. To that end, in various implementations, the image processing architecture **556** includes instructions and/or logic therefor, and heuristics and metadata therefor.

[0055] In some implementations, the compositor **558** is configured to composite the rendered A/V content and/or XR content with the processed image stream of the physical environment **105** from the image processing architecture **556** to produce rendered image frames of the XR environment **128** for display. The compositor **558** is described in more detail below with reference to FIG. **5**. To that end, in various implementations, the compositor **558** includes instructions and/or logic therefor, and heuristics and metadata therefor.

[0056] Although the data obtainer **242**, the mapper and locator engine **244**, the data transmitter **246**, the privacy architecture **408**, the motion state estimator **410**, the eye tracking engine **412**, the head/body pose tracking engine **414**, the characterization engine **416**, the context analyzer **460**, the muscle strain engine **463**, the interactive stretching engine **468A**, the API **468B**, the content selector **522**, the content manager **530**, and the rendering engine **550** are

shown as residing on a single device (e.g., the controller **110**), it should be understood that in other implementations, any combination of the data obtainer **242**, the mapper and locator engine **244**, the data transmitter **246**, the privacy architecture **408**, the motion state estimator **410**, the eye tracking engine **412**, the head/body pose tracking engine **414**, the characterization engine **416**, the context analyzer **460**, the muscle strain engine **463**, the interactive stretching engine **468A**, the API **468B**, the content selector **522**, the content manager **530**, and the rendering engine **550** may be located in separate computing devices.

[0057] In some implementations, the functions and/or components of the controller **110** are combined with or provided by the electronic device **120** shown below in FIG. 3. Moreover, FIG. 2 is intended more as a functional description of the various features which may be present in a particular implementation as opposed to a structural schematic of the implementations described herein. As recognized by those of ordinary skill in the art, items shown separately could be combined and some items could be separated. For example, some functional modules shown separately in FIG. 2 could be implemented in a single module and the various functions of single functional blocks could be implemented by one or more functional blocks in various implementations. The actual number of modules and the division of particular functions and how features are allocated among them will vary from one implementation to another and, in some implementations, depends in part on the particular combination of hardware, software, and/or firmware chosen for a particular implementation.

[0058] FIG. 3 is a block diagram of an example of the electronic device **120** (e.g., a mobile phone, tablet, laptop, near-eye system, wearable computing device, or the like) in accordance with some implementations. While certain specific features are illustrated, those skilled in the art will appreciate from the present disclosure that various other features have not been illustrated for the sake of brevity, and so as not to obscure more pertinent aspects of the implementations disclosed herein. To that end, as a non-limiting example, in some implementations, the electronic device **120** includes one or more processing units **302** (e.g., microprocessors, ASICs, FPGAs, GPUs, CPUs, processing cores, and/or the like), one or more input/output (I/O) devices and sensors **306**, one or more communication interfaces **308** (e.g., USB, IEEE 802.3x, IEEE 802.11x, IEEE 802.16x, GSM, CDMA, TDMA, GPS, IR, BLUETOOTH, ZIGBEE, and/or the like type interface), one or more programming (e.g., I/O) interfaces **310**, one or more displays **312**, an image capture device **370** (e.g., one or more optional interior-facing and/or exterior-facing image sensors), a memory **320**, and one or more communication buses **304** for interconnecting these and various other components.

[0059] In some implementations, the one or more communication buses **304** include circuitry that interconnects and controls communications between system components. In some implementations, the one or more I/O devices and sensors **306** include at least one of an inertial measurement unit (IMU), an accelerometer, a gyroscope, a magnetometer, a thermometer, one or more physiological sensors (e.g., blood pressure monitor, heart rate monitor, blood oximetry monitor, blood glucose monitor, etc.), one or more microphones, one or more speakers, a haptics engine, a heating and/or cooling unit, a skin shear engine, one or more depth sensors (e.g., structured light, time-of-flight, LiDAR, or the

like), a localization and mapping engine, an eye tracking engine, a head/body pose tracking engine, a hand/limb/finger/extremity tracking engine, a camera pose tracking engine, and/or the like.

[0060] In some implementations, the one or more displays **312** are configured to present the XR environment to the user. In some implementations, the one or more displays **312** are also configured to present flat video content to the user (e.g., a 2-dimensional or “flat” AVI, FLV, WMV, MOV, MP4, or the like file associated with a TV episode or a movie, or live video pass-through of the physical environment **105**). In some implementations, the one or more displays **312** correspond to touchscreen displays (e.g., similar to the display **122** in FIG. 1). In some implementations, the one or more displays **312** correspond to holographic, digital light processing (DLP), liquid-crystal display (LCD), liquid-crystal on silicon (LCoS), organic light-emitting field-effect transitory (OLET), organic light-emitting diode (OLED), surface-conduction electron-emitter display (SED), field-emission display (FED), quantum-dot light-emitting diode (QD-LED), micro-electro-mechanical system (MEMS), and/or the like display types. In some implementations, the one or more displays **312** correspond to diffractive, reflective, polarized, holographic, etc. waveguide displays. For example, the electronic device **120** includes a single display such as the display **122**. In another example, the electronic device **120** includes a display for each eye of the user. In some implementations, the one or more displays **312** are capable of presenting AR and VR content. In some implementations, the one or more displays **312** are capable of presenting AR or VR content.

[0061] In some implementations, the image capture device **370** correspond to one or more RGB cameras (e.g., with a complementary metal-oxide-semiconductor (CMOS) image sensor or a charge-coupled device (CCD) image sensor), IR image sensors, event-based cameras, and/or the like. In some implementations, the image capture device **370** includes a lens assembly, a photodiode, and a front-end architecture. In some implementations, the image capture device **370** includes exterior-facing and/or interior-facing image sensors.

[0062] The memory **320** includes high-speed random-access memory such as DRAM, SRAM, DDR RAM, or other random-access solid-state memory devices. In some implementations, the memory **320** includes non-volatile memory such as one or more magnetic disk storage devices, optical disk storage devices, flash memory devices, or other non-volatile solid-state storage devices. The memory **320** optionally includes one or more storage devices remotely located from the one or more processing units **302**. The memory **320** comprises a non-transitory computer readable storage medium. In some implementations, the memory **320** or the non-transitory computer readable storage medium of the memory **320** stores the following programs, modules and data structures, or a subset thereof including an optional operating system **330** and a presentation engine **340**.

[0063] The operating system **330** includes procedures for handling various basic system services and for performing hardware dependent tasks. In some implementations, the presentation engine **340** is configured to present media items and/or XR content to the user via the one or more displays **312**. To that end, in various implementations, the presentation engine **340** includes a data obtainer **342**, an interaction handler **520**, a presenter **560**, and a data transmitter **350**.

[0064] In some implementations, the data obtainer 342 is configured to obtain data (e.g., presentation data such as rendered image frames associated with the user interface or the XR environment, input data, user interaction data, head tracking information, camera pose tracking information, eye tracking information, hand/limb/finger/extremity tracking information, sensor data, location data, etc.) from at least one of the I/O devices and sensors 306 of the electronic device 120, the controller 110, and the remote input devices. To that end, in various implementations, the data obtainer 342 includes instructions and/or logic therefor, and heuristics and metadata therefor.

[0065] In some implementations, the interaction handler 520 is configured to detect user interactions (e.g., gestural inputs detected via hand/extremity tracking, eye gaze inputs detected via eye tracking, voice commands, etc.) with the presented A/V content and/or XR content. To that end, in various implementations, the interaction handler 520 includes instructions and/or logic therefor, and heuristics and metadata therefor.

[0066] In some implementations, the presenter 560 is configured to present and update A/V content and/or XR content (e.g., the rendered image frames associated with the user interface or the XR environment 128 including the VA(s), the XR content, one or more UI elements associated with the XR content, and/or the like) via the one or more displays 312. To that end, in various implementations, the presenter 560 includes instructions and/or logic therefor, and heuristics and metadata therefor.

[0067] In some implementations, the data transmitter 350 is configured to transmit data (e.g., presentation data, location data, user interaction data, head tracking information, camera pose tracking information, eye tracking information, hand/limb/finger/extremity tracking information, etc.) to at least the controller 110. To that end, in various implementations, the data transmitter 350 includes instructions and/or logic therefor, and heuristics and metadata therefor.

[0068] Although the data obtainer 342, the interaction handler 520, the presenter 560, and the data transmitter 350 are shown as residing on a single device (e.g., the electronic device 120), it should be understood that in other implementations, any combination of the data obtainer 342, the interaction handler 520, the presenter 560, and the data transmitter 350 may be located in separate computing devices.

[0069] Moreover, FIG. 3 is intended more as a functional description of the various features which may be present in a particular implementation as opposed to a structural schematic of the implementations described herein. As recognized by those of ordinary skill in the art, items shown separately could be combined and some items could be separated. For example, some functional modules shown separately in FIG. 3 could be implemented in a single module and the various functions of single functional blocks could be implemented by one or more functional blocks in various implementations. The actual number of modules and the division of particular functions and how features are allocated among them will vary from one implementation to another and, in some implementations, depends in part on the particular combination of hardware, software, and/or firmware chosen for a particular implementation.

[0070] FIG. 4A is a block diagram of a first portion 400A of an example data processing architecture in accordance with some implementations. While pertinent features are

shown, those of ordinary skill in the art will appreciate from the present disclosure that various other features have not been illustrated for the sake of brevity and so as not to obscure more pertinent aspects of the example implementations disclosed herein. To that end, as a non-limiting example, the first portion 400A of the data processing architecture is included in a computing system such as the controller 110 shown in FIGS. 1 and 2; the electronic device 120 shown in FIGS. 1 and 3; and/or a suitable combination thereof.

[0071] As shown in FIG. 4A, one or more local sensors 402 of the controller 110, the electronic device 120, and/or a combination thereof obtain local sensor data 403 associated with the physical environment 105. For example, the local sensor data 403 includes images or a stream thereof of the physical environment 105, simultaneous location and mapping (SLAM) information for the physical environment 105 and the location of the electronic device 120 or the user 150 relative to the physical environment 105, ambient lighting information for the physical environment 105, ambient audio information for the physical environment 105, acoustic information for the physical environment 105, dimensional information for the physical environment 105, semantic labels for objects within the physical environment 105, and/or the like. In some implementations, the local sensor data 403 includes un-processed or post-processed information.

[0072] Similarly, as shown in FIG. 4A, one or more remote sensors 404 associated with the optional remote input devices within the physical environment 105 obtain remote sensor data 405 associated with the physical environment 105. For example, the remote sensor data 405 includes images or a stream thereof of the physical environment 105, SLAM information for the physical environment 105 and the location of the electronic device 120 or the user 150 relative to the physical environment 105, ambient lighting information for the physical environment 105, ambient audio information for the physical environment 105, acoustic information for the physical environment 105, dimensional information for the physical environment 105, semantic labels for objects within the physical environment 105, and/or the like. In some implementations, the remote sensor data 405 includes un-processed or post-processed information.

[0073] According to some implementations, the privacy architecture 408 ingests the local sensor data 403 and the remote sensor data 405. In some implementations, the privacy architecture 408 includes one or more privacy filters associated with user information and/or identifying information. In some implementations, the privacy architecture 408 includes an opt-in feature where the electronic device 120 informs the user 150 as to what user information and/or identifying information is being monitored and how the user information and/or the identifying information will be used. In some implementations, the privacy architecture 408 selectively prevents and/or limits the data processing architecture 400A/400B/400C or portions thereof from obtaining and/or transmitting the user information. To this end, the privacy architecture 408 receives user preferences and/or selections from the user 150 in response to prompting the user 150 for the same. In some implementations, the privacy architecture 408 prevents the data processing architecture 400A/400B/400C from obtaining and/or transmitting the user information unless and until the privacy architecture



**408** obtains informed consent from the user **150**. In some implementations, the privacy architecture **408** anonymizes (e.g., scrambles, obscures, encrypts, and/or the like) certain types of user information. For example, the privacy architecture **408** receives user inputs designating which types of user information the privacy architecture **408** anonymizes. As another example, the privacy architecture **408** anonymizes certain types of user information likely to include sensitive and/or identifying information, independent of user designation (e.g., automatically).

[0074] According to some implementations, the motion state estimator **410** obtains the local sensor data **403** and the remote sensor data **405** after it has been subjected to the privacy architecture **408**. In some implementations, the motion state estimator **410** obtains (e.g., receives, retrieves, or determines/generates) a motion state vector **411** based on the input data and updates the motion state vector **411** over time.

[0075] FIG. 4B shows an example data structure for the motion state vector **411** in accordance with some implementations. As shown in FIG. 4B, the motion state vector **411** may correspond to an N-tuple characterization vector or characterization tensor that includes a timestamp **421** (e.g., the most recent time the motion state vector **411** was updated), a motion state descriptor **422** for the electronic device **120** (e.g., stationary, in-motion, running, walking, cycling, driving or riding in a car, driving or riding in a boat, driving or riding in a bus, riding in a train, riding in a plane, or the like), translational movement values **424** associated with the electronic device **120** (e.g., a heading, a displacement value, a velocity value, an acceleration value, a jerk value, etc.), angular movement values **426** associated with the electronic device **120** (e.g., an angular displacement value, an angular velocity value, an angular acceleration value, an angular jerk value, and/or the like for each of the pitch, roll, and yaw dimensions), and/or miscellaneous information **428**. One of ordinary skill in the art will appreciate that the data structure for the motion state vector **411** in FIG. 4B is merely an example that may include different information portions in various other implementations and be structured in myriad ways in various other implementations.

[0076] According to some implementations, the eye tracking engine **412** obtains the local sensor data **403** and the remote sensor data **405** after it has been subjected to the privacy architecture **408**. In some implementations, the eye tracking engine **412** obtains (e.g., receives, retrieves, or determines/generates) an eye tracking vector **413** based on the input data and updates the eye tracking vector **413** over time.

[0077] FIG. 4B shows an example data structure for the eye tracking vector **413** in accordance with some implementations. As shown in FIG. 4B, the eye tracking vector **413** may correspond to an N-tuple characterization vector or characterization tensor that includes a timestamp **431** (e.g., the most recent time the eye tracking vector **413** was updated), one or more angular values **432** for a current gaze direction (e.g., roll, pitch, and yaw values), one or more translational values **434** for the current gaze direction (e.g., x, y, and z values relative to the physical environment **105**, the world-at-large, and/or the like), and/or miscellaneous information **436**. One of ordinary skill in the art will appreciate that the data structure for the eye tracking vector **413** in FIG. 4B is merely an example that may include

different information portions in various other implementations and be structured in myriad ways in various other implementations.

[0078] For example, the gaze direction indicates a point (e.g., associated with x, y, and z coordinates relative to the physical environment **105** or the world-at-large), a physical object, or a region of interest (ROI) in the physical environment **105** at which the user **150** is currently looking. As another example, the gaze direction indicates a point (e.g., associated with x, y, and z coordinates relative to the XR environment **128**), an XR object, or a region of interest (ROI) in the XR environment **128** at which the user **150** is currently looking.

[0079] According to some implementations, the head/body pose tracking engine **414** obtains the local sensor data **403** and the remote sensor data **405** after it has been subjected to the privacy architecture **408**. In some implementations, the head/body pose tracking engine **414** obtains (e.g., receives, retrieves, or determines/generates) a pose characterization vector **415** based on the input data and updates the pose characterization vector **415** over time.

[0080] FIG. 4B shows an example data structure for the pose characterization vector **415** in accordance with some implementations. As shown in FIG. 4B, the pose characterization vector **415** may correspond to an N-tuple characterization vector or characterization tensor that includes a timestamp **441** (e.g., the most recent time the pose characterization vector **415** was updated), a head pose descriptor **442A** (e.g., upward, downward, neutral, etc.), translational values for the head pose **442B**, rotational values for the head pose **442C**, a body pose descriptor **444A** (e.g., standing, sitting, prone, etc.), translational values for body sections/extremities/limbs/joints **444B**, rotational values for the body sections/extremities/limbs/joints **444C**, and/or miscellaneous information **446**. In some implementations, the pose characterization vector **415** also includes information associated with finger/hand/extremity tracking. One of ordinary skill in the art will appreciate that the data structure for the pose characterization vector **415** in FIG. 4B is merely an example that may include different information portions in various other implementations and be structured in myriad ways in various other implementations. According to some implementations, the motion state vector **411**, the eye tracking vector **413** and the pose characterization vector **415** are collectively referred to as an input vector **417**.

[0081] According to some implementations, the characterization engine **416** obtains the motion state vector **411**, the eye tracking vector **413** and the pose characterization vector **415**. In some implementations, the characterization engine **416** obtains (e.g., receives, retrieves, or determines/generates) the characterization vector **419** based on the motion state vector **411**, the eye tracking vector **413**, and the pose characterization vector **415**.

[0082] FIG. 4B shows an example data structure for the characterization vector **419** in accordance with some implementations. As shown in FIG. 4B, the characterization vector **419** may correspond to an N-tuple characterization vector or characterization tensor that includes a timestamp **451** (e.g., the most recent time the characterization vector **419** was updated), motion state information **452** (e.g., the motion state descriptor **422**), gaze direction information **454** (e.g., a function of the one or more angular values **432** and the one or more translational values **434** within the eye tracking vector **413**), head pose information **456A** (e.g., a

function of the head pose descriptor **442A** within the pose characterization vector **415**), body pose information **456B** (e.g., a function of the body pose descriptor **444A** within the pose characterization vector **415**), extremity tracking information **456C** (e.g., a function of the body pose descriptor **444A** within the pose characterization vector **415** that is associated with extremities of the user **150** that are being tracked by the controller **110**, the electronic device **120**, and/or a combination thereof), location information **458** (e.g., a household location such as a kitchen or living room, a vehicular location such as an automobile, plane, etc., and/or the like), and/or miscellaneous information **459**.

[0083] FIG. 4C is a block diagram of a second portion **400B** of the example data processing architecture in accordance with some implementations. While pertinent features are shown, those of ordinary skill in the art will appreciate from the present disclosure that various other features have not been illustrated for the sake of brevity and so as not to obscure more pertinent aspects of the example implementations disclosed herein. To that end, as a non-limiting example, the second portion **400B** of the data processing architecture is included in a computing system such as the controller **110** shown in FIGS. 1 and 2; the electronic device **120** shown in FIGS. 1 and 3; and/or a suitable combination thereof. FIG. 4C is similar to and adapted from FIG. 4A. Therefore, similar reference numbers are used in FIGS. 4A and 4C. As such, only the differences between FIGS. 4A and 4C will be described below for the sake of brevity.

[0084] According to some implementations, the context analyzer **460** obtains the motion state vector **411** from the motion state estimator **410**. As shown in FIG. 4C, the context analyzer **460** also obtains the local sensor data **403** and the remote sensor data **405** after being subjected to the privacy architecture **408**.

[0085] In some implementations, the context analyzer **460** obtains (e.g., receives, retrieves, or determines/generates) a context information vector **470** based on the input data and updates the context information vector **470** over time. FIG. 4C shows an example data structure for the context information vector **470** in accordance with some implementations. As shown in FIG. 4C, the context information vector **470** may correspond to an N-tuple characterization vector or characterization tensor that includes: a timestamp **471** (e.g., the most recent time the context information vector **470** was updated); environment state information **472** associated with a current state of the physical environment **105** (e.g., ambient temperature information, ambient humidity information, ambient lighting information, ambient audio information, semantic labels for physical objects within the physical environment **105**, locations for physical objects within the physical environment **105**, etc.); device state information **474** associated with a current state of the controller **110**, the electronic device **120**, or a combination thereof, or the like (e.g., current foreground applications, current background applications, power/charge remaining, device temperature metrics, resource consumption metrics (e.g., CPU, RAM, storage, network I/O, etc.), etc.); user state information **476** associated with a current state of the user **150** (e.g., the characterization vector **419**, physiological information associated with the user **150**, the motion state descriptor **42**, etc.); and miscellaneous information **478**. One of ordinary skill in the art will appreciate that the data structure for the context information vector **470** in FIG. 4C is merely an example that

may include different information portions in various other implementations and be structured in myriad ways in various other implementations.

[0086] According to some implementations, the head/body/neck mechanics engine **462** obtains (e.g., receives, retrieves, or determines/generates) displacement, velocity, acceleration, jerk, torque, etc. values for the head/body/neck of the user **150** based on changes to the pose characterization vector **415**. In some implementations, the strain analyzer **464** determines current strain information **480** for one or more muscles or muscles groups based on: the displacement, velocity, acceleration, jerk, torque, etc. values for the head/body/neck of the user **150** from the head/body/neck mechanics engine **462**; historical information **466**; and the context information vector **470**. In some implementations, the strain analyzer **464** determines the current strain information **480** based on strain increase logic **465A** and/or strain decrease logic **465B**. In some implementations, the historical information **466** corresponds to local or remote storage repository, including: strain information for one or more previous time periods on an overall basis, individual muscle or muscle group/region basis, etc.; context information for one or more previous time periods; and/or displacement, velocity, acceleration, jerk, torque, etc. values for the head/body/neck of the user **150** for one or more previous time periods.

[0087] FIG. 4D shows an example data structure for the current strain information **480** in accordance with some implementations. As shown in FIG. 4C, the current strain information **480** may correspond to an N-tuple characterization vector or characterization tensor that includes: a timestamp **481**; muscle information **482A** associated with a first muscle or muscle group/region; muscle information **482B** associated with a second muscle or muscle group/region; muscle information **482C** associated with a third muscle or muscle group/region; muscle information **482D** associated with a fourth muscle or muscle group/region; muscle information **482E** associated with a fifth muscle or muscle group/region; accumulated strain information **486** associated with a function of one or more of the muscle information **482A-482E**; and miscellaneous information **488**. One of ordinary skill in the art will appreciate that the data structure for current strain information **480** in FIG. 4D is merely an example that may include different information portions in various other implementations and be structured in myriad ways in various other implementations.

[0088] As shown in FIG. 4D, the muscle information **482A** for the first muscle or muscle group/region includes: a muscle identifier **484A** for the first muscle or muscle group/region (e.g., a unique identifier, a label, a name, or the like for the first muscle or muscle group/region); a current muscle strain value **485A** for the first muscle or muscle group/region; a pointer to historical muscle strain information **486A** for the first muscle or muscle group/region within the historical information **466**; and miscellaneous information **487A** associated with the first muscle or muscle group/region.

[0089] As shown in FIG. 4D, for example, the muscle strain engine **463** determines current muscle strain values **485A**, **485B**, **485C**, **485D**, and **485E** for muscles or muscle groups/regions **484A**, **484B**, **484C**, **484D**, and **484E**, respectively, of the user **150**. Furthermore, the muscle strain engine **463** updates (increases or decreases) the muscle strain values **485A**, **485B**, **485C**, **485D**, and **485E** over time based on

rotational and/or translational movement of the user **150** that triggers the strain increase logic **465A** and/or the strain decrease logic **465B**.

[0090] As shown in FIG. 4D, the accumulated strain information **486**: a current accumulated strain value **489A** associated with a function of one or more of the current muscle strain values **485A**, **485B**, **485C**, **485D**, and **485E** for muscles or muscle groups/regions **484A**, **484B**, **484C**, **484D**, and **484E**, respectively, of the user **150**; a pointer to historical accumulated strain information **489B** within the historical information **466**; and miscellaneous information **489C** associated with the accumulated strain information. As shown in FIG. 4D, for example, the muscle strain engine **463** also determines a current accumulated strain value **489A** and updates (increases or decreases) the current accumulated strain value **489A** over time based on rotational and/or translational movement of the user **150** that triggers the strain increase logic **465A** and/or the strain decrease logic **465B**. As such, according to some implementations, the muscle strain engine **463** tracks strain values on an individual muscle or muscle group/region basis (e.g., the muscle information **482A-482E**) as well as an overall strain value (e.g., the current accumulated strain information **486**).

[0091] FIG. 4E is a block diagram of a third portion **400C** of the example data processing architecture in accordance with some implementations. While pertinent features are shown, those of ordinary skill in the art will appreciate from the present disclosure that various other features have not been illustrated for the sake of brevity and so as not to obscure more pertinent aspects of the example implementations disclosed herein. To that end, as a non-limiting example, the third portion **400C** of the example data processing architecture is included in a computing system such as the controller **110** shown in FIGS. 1 and 2; the electronic device **120** shown in FIGS. 1 and 3; and/or a suitable combination thereof. FIG. 4E is similar to and adapted from FIGS. 4A and 4C. Therefore, similar reference numbers are used in FIGS. 4A and 4C. As such, only the differences between FIGS. 4A, 4C, and 4E will be described below for the sake of brevity.

[0092] As described above with respect to FIG. 4C, the muscle strain engine **463** determines a current strain information **480**. As illustrated in FIG. 4E, the current strain information **480** is provided to a multiplexer (Mux) **467**. In turn, the current strain information **480** is provided to at least one of an interactive stretching engine **468A** and an application programming interface (API) **468B**. According to some implementations, the interactive stretching engine **468A** determines/generates stretching session visualizations **469** (e.g., the first plurality of visual cues **614A**, **614B**, **614C** and the directional indicator **612A** in FIG. 6A) based on the characterization vector **419** (e.g., including head and/or body pose information), user profile information **473**, and/or the context information vector **470**. The stretching session visualizations **469** are described in greater detail below with reference to FIGS. 6A-6F.

[0093] In some implementations, the interactive stretching engine **468A** generates the stretching session visualizations **469** in response to detecting a user input to initiate a guided stretching session (e.g., a touch input, a hand/extremity tracking input, an eye tracking input, a voice command, or the like). In some implementations, the interactive stretching engine **468A** generates the stretching session visualizations **469** in response to detecting an accumulated strain value

greater than a threshold value. In some implementations, the threshold value corresponds to a deterministic or non-deterministic value. According to some implementations, the threshold value corresponds to the fourth posture awareness threshold mentioned in provisional patent application No xxx,xxx, filed on xxx (Attorney Docket Number 27753-50497US1), which is incorporated by reference in its entirety.

[0094] According to some implementations, the API **468B** provides access to the current strain information **480** to at least one of: the operating system of the controller **110**, the electronic device **120**, or a combination thereof; third-party programs or applications; and/or the like. As such, the current strain information **480** may be used in various downstream processes.

[0095] FIG. 5 is a block diagram of an example content delivery architecture **500** in accordance with some implementations. While pertinent features are shown, those of ordinary skill in the art will appreciate from the present disclosure that various other features have not been illustrated for the sake of brevity and so as not to obscure more pertinent aspects of the example implementations disclosed herein. To that end, as a non-limiting example, the content delivery architecture is included in a computing system such as the controller **110** shown in FIGS. 1 and 2; the electronic device **120** shown in FIGS. 1 and 3; and/or a suitable combination thereof.

[0096] According to some implementations, the interaction handler **520** obtains (e.g., receives, retrieves, or detects) one or more user inputs **521** provided by the user **150** that are associated with selecting A/V content, one or more VAs, and/or XR content for presentation. For example, the one or more user inputs **521** correspond to a gestural input selecting XR content from a UI menu detected via hand/extremity tracking, an eye gaze input selecting XR content from the UI menu detected via eye tracking, a voice command selecting XR content from the UI menu detected via a microphone, and/or the like. In some implementations, the content selector **522** selects XR content **527** from the content library **525** based on one or more user inputs **521** (e.g., a voice command, a selection from a menu of XR content items, and/or the like).

[0097] In various implementations, the content manager **530** manages and updates the layout, setup, structure, and/or the like for the XR environment **128**, including one or more of VAs, XR content, one or more UI elements associated with the XR content, and/or the like, based on the characterization vector **419**, (optionally) the user inputs **521**, and/or the like. To that end, the content manager **530** includes the frame buffer **532**, the content updater **534**, and the feedback engine **536**.

[0098] In some implementations, the frame buffer **532** includes XR content, a rendered image frame, and/or the like for one or more past instances and/or frames. In some implementations, the content updater **534** modifies the XR environment **128** over time based on the characterization vector **419**, the stretching session visualizations **469** (e.g., the first plurality of visual cues **614A**, **614B**, **614C** and the directional indicator **612A** in FIG. 6A), the user inputs **521** associated with modifying and/or manipulating the XR content or VA(s), translational or rotational movement of objects within the physical environment **105**, translational or rotational movement of the electronic device **120** (or the user **150**), and/or the like. In some implementations, the

feedback engine **536** generates sensory feedback (e.g., visual feedback such as text or lighting changes, audio feedback, haptic feedback, etc.) associated with the XR environment **128**.

[0099] According to some implementations, the pose determiner **552** determines a current camera pose of the electronic device **120** and/or the user **150** relative to the XR environment **128** and/or the physical environment **105** based at least in part on the pose characterization vector **415**. In some implementations, the renderer **554** renders the VA(s), the XR content **527**, one or more UI elements associated with the XR content, and/or the like according to the current camera pose relative thereto.

[0100] According to some implementations, the optional image processing architecture **556** obtains an image stream from an image capture device **370** including one or more images of the physical environment **105** from the current camera pose of the electronic device **120** and/or the user **150**. In some implementations, the image processing architecture **556** also performs one or more image processing operations on the image stream such as warping, color correction, gamma correction, sharpening, noise reduction, white balance, and/or the like. In some implementations, the optional compositor **558** composites the rendered XR content with the processed image stream of the physical environment **105** from the image processing architecture **556** to produce rendered image frames of the XR environment **128**. In various implementations, the presenter **560** presents the rendered image frames of the XR environment **128** to the user **150** via the one or more displays **312**. One of ordinary skill in the art will appreciate that the optional image processing architecture **556** and the optional compositor **558** may not be applicable for fully virtual environments (or optical see-through scenarios).

[0101] FIGS. 6A-6F illustrate a plurality of three-dimensional (3D) environments associated with a guided stretching session in accordance with some implementations. While certain specific features are illustrated, those skilled in the art will appreciate from the present disclosure that various other features have not been illustrated for the sake of brevity, and so as not to obscure more pertinent aspects of the implementations disclosed herein. To that end, as a non-limiting example, the plurality of interfaces is rendered and presented by a computing system such as the controller **110** shown in FIGS. 1 and 2; the electronic device **120** shown in FIGS. 1 and 3; and/or a suitable combination thereof.

[0102] In some implementations, the plurality of 3D environments in FIGS. 6A-6F correspond to the XR environment **128** shown in FIG. 1 (e.g., a 3D or volumetric user interface). As such, according to some implementations, the electronic device **120** presents the plurality of 3D environments to the user **150** while the user **150** is physically present within a physical environment, which is currently within the FOV **111** of an exterior-facing image sensor of the electronic device **120** (e.g., as shown in FIG. 1). In other words, in some implementations, the electronic device **120** is configured to present XR content (e.g., virtual content) and to enable optical see-through or video pass-through of at least a portion of the physical environment on the display **122**. For example, the electronic device **120** corresponds to a mobile phone, tablet, laptop, near-eye system, wearable computing device, or the like. As such, in some implemen-

tations, the user **150** holds the electronic device **120** in their hand(s) similar to the operating environment **100** in FIG. 1.

[0103] As shown in FIG. 6A, the electronic device **120** presents, on the display **122**, a 3D environment **602A** associated with the guided stretching session at time **T1**. Continuing with this example, while presenting the 3D environment **602A**, the electronic device **120** presents a notification **610** welcoming the user to the guided stretching session. In some implementations, the notification **610** (and also the notifications **620**, **630**, **640**, **650**, and **660** in FIGS. 6B-6F, respectively) corresponds to two-dimensional (2D) content overlaid on the 3D environment. For example, the electronic device **120** presents the notification **610** at a default location or in a default manner in FIG. 6A (e.g., a pop-up notification centered within the display **122**, a banner notification adjacent to the top edge of the display **122**, or the like). In some implementations, the notification **610** (and also the notifications **620**, **630**, **640**, **650**, and **660** in FIGS. 6B-6F, respectively) corresponds to 3D or volumetric content overlaid on or composited with the 3D environment.

[0104] As shown in FIG. 6A, the 3D environment **602A** includes a first plurality of visual cues (e.g., detents) **614A**, **614B**, **614C** for a first portion of a guided stretching session (e.g., a 90° rotational yaw movement of the user's head). For example, the first plurality of visual cues **614A**, **614B**, **614C** is evenly spaced along a 90° arc relative to the user's current location. For example, if the first plurality of visual cues includes three visual cues, the electronic device **120** presents the three visual cues on the 90° arc with 30° between the starting, neutral head pose and each of the three visual cues. In some implementations, the first plurality of visual cues **614A**, **614B**, **614C** have similar sizes but appear different based on their depth and/or distance relative to the user's current location.

[0105] One of ordinary skill in the art will appreciate that the usage of rotational yaw movement for the guided stretching session in FIGS. 6A-6F is merely an example that may be replaced with one or more of rotational yaw, pitch, or roll movement. According to some implementations, with respect to yaw rotational movement, the first plurality of visual cues remains within the user's field of view. According to some implementations, with respect to pitch rotational movement, the first plurality of visual cues remains within the user's field of view. In implementations, with respect to pitch rotational movement, the first plurality of visual cues remains horizontally aligned (e.g., perpendicular to an alignment vector such as gravity).

[0106] According to some implementations, the electronic device **120** determines locations for the first plurality of visual cues **614A**, **614B**, **614C** within the 3D environment **602A** based on user profile information and the head pose information associated with the user. For example, the user profile information corresponds to visual acuity metrics for the user, user height, user preferences, and/or the like. For example, the head pose information at least includes three-degrees of freedom (3DOF) rotational values determined by, for example, the head/body pose tracking engine **414** described above with reference to FIGS. 2, 4A, and 4C.

[0107] In some implementations, the electronic device **120** determines the locations for the first plurality of visual cues **614A**, **614B**, **614C** to suit the user's height and/or sightline (e.g., in front of and level with the user's sightline). In some implementations, the electronic device **120** determines the locations for the first plurality of visual cues **614A**, **614B**,

**614C** so as not to occlude or collide with objects within the 3D environment. In some implementations, the electronic device **120** determines the locations for the first plurality of visual cues **614A**, **614B**, **614C** based on user flexibility/range of motion (taking user history into account) or user preferences. For example, the angular range associated with the first plurality of visual cues **614A**, **614B**, **614C** may be limited by the user's maximum range of observed motion, a present maximum range of motion provided by the user, or a historical range of motion associated with the user.

[0108] As shown in FIG. 6A, the 3D environment **602A** also includes a directional indicator **612A** associated with a first stretch direction for the first portion of a guided stretching session (e.g., left-to-right or clockwise rotational movement). In one example, the electronic device **120** presents the directional indicator **612A** at a deterministic (default or predefined) origin location when starting the guided stretching session. In another example, the electronic device **120** presents the directional indicator **612A** at a non-deterministic origin location based on the user's current height and/or sightline. As shown in FIG. 6A, the directional indicator **612A** includes a visual representation (e.g., the number "3") associated with a count of the first plurality of visual cues that have not been completed for the first portion of the guided stretching session at time T1.

[0109] As shown in FIG. 6B, the electronic device **120** presents, on the display **122**, a 3D environment **602B** associated with the first portion of the guided stretching session at time T2. Continuing with this example, while presenting the 3D environment **602B**, the electronic device **120** presents a notification **620** instructing the user to follow the visual cues for the first portion of the guided stretching session by rotating their head from left-to-right by 90°. As shown in FIG. 6B, while presenting the 3D environment **602B**, the electronic device **120** detects a change **622** to the head pose information associated with the user **150** such as a 30° clockwise rotational yaw movement. In some implementations, the electronic device **120** presents a representation of the user **150** within the 3D environment that tracks the head pose and/or body pose changes associated with the user **150**. As such, in this example, the electronic device **120** updates the representation of the user **150** based on the change **622** to the head pose information associated with the user **150**. For example, the representation of the user **150** corresponds to an avatar with a face and/or body modeled on the user **150**.

[0110] In response to detecting the change **622** to the head pose information associated with the user **150** in FIG. 6B, the electronic device **120** presents, on the display **122**, a 3D environment **602C** associated with the first portion of the guided stretching session at time T3 in FIG. 6C. In response to detecting the change **622** to the head pose information associated with the user **150** in FIG. 6B and in accordance with a determination that the change **622** to the head pose information associated with the user **150** satisfies a criterion associated with the visual cue **614A**, the electronic device **120** presents a notification **630** within the 3D environment **602C** (e.g., visual feedback) indicating that the visual cue **614A** has been completed for the first portion of the guided stretching session and to continue with the same motion. In some implementations, the criterion associated with a visual cue is satisfied when the location for the directional indicator coincides with the location for the visual cue.

[0111] In some implementations, the visual feedback is replaced with or supplemented by haptic feedback, audio feedback, and/or the like. According to some implementations, each of the visual cues, save for the final visual cue, is associated with similar haptic and/or audio feedback, and the final visual cue is associated with distinct haptic and/or audio feedback with a greater intensity, length, pitch, frequency, volume, and/or the like. According to some implementations, audio feedback may be spatialized to coincide with the location of the visual cue or at a distance in a direction of the stretch that satisfies the visual cue. In one example, with reference to yaw rotational movement, the audio feedback may be spatialized at a distance and at the angle such that the user moves their head side-to-side about the y axis (e.g., based on the user's head and sightline) to satisfy the current portion of the guided stretching session. In another example, with reference to roll rotational movement for the guided stretching session, the audio feedback may be spatialized relative to a location above the user's head and at an angle such that the user tilts their head about the z axis (e.g., based on the user's head and sightline) to satisfy the current portion of the guided stretching session. In yet another example, with reference to pitch rotational movement, the audio feedback may be spatialized at a location in front of the user and at an angle such that the user moves their head about the x axis (e.g., based on the user's head and sightline) to satisfy the current portion of the guided stretching session.

[0112] As shown in FIG. 6C, the 3D environment **602C** includes the visual cues **614B** and **614C** and the directional indicator **612B** at the former location of the visual cue **614A**. As shown in FIG. 6C, the directional indicator **612B** includes a visual representation (e.g., the number "2") associated with a count of the first plurality of visual cues that have not been completed for the first portion of the guided stretching session at time T3. As shown in FIG. 6C, while presenting the 3D environment **602C**, the electronic device **120** detects a change **632** to the head pose information associated with the user **150** such as another 30° clockwise rotational yaw movement.

[0113] In response to detecting the change **632** to the head pose information associated with the user **150** in FIG. 6C, the electronic device **120** presents, on the display **122**, a 3D environment **602D** associated with the first portion of the guided stretching session at time T4 in FIG. 6D. In response to detecting the change **632** to the head pose information associated with the user **150** in FIG. 6C and in accordance with a determination that the change **632** to the head pose information associated with the user **150** satisfies a criterion associated with the visual cue **614B**, the electronic device **120** presents a notification **640** within the 3D environment **602D** (e.g., visual feedback) indicating that the visual cue **614B** has been completed for the first portion of the guided stretching session and to continue with the same motion. In some implementations, the visual feedback is replaced with or supplemented by haptic feedback, audio feedback, spatial audio feedback, and/or the like as described above with reference to FIG. 6B.

[0114] As shown in FIG. 6D, the 3D environment **602D** includes the visual cue **614C** and the directional indicator **612C** at the former location of visual cue **614B**. As shown in FIG. 6D, the directional indicator **612C** includes a visual representation (e.g., the number "1") associated with a count of the first plurality of visual cues that have not been

completed for the first portion of the guided stretching session at time T4. As shown in FIG. 6D, while presenting the 3D environment 602D, the electronic device 120 detects a change 642 to the head pose information associated with the user 150 such as yet another 30° clockwise rotational yaw movement. In one example, the changes 622, 632, and 642 are included in one continuous 90° clockwise rotational yaw movement. In another example, the changes 622, 632, and 642 correspond to separate, sequential 30° clockwise rotational yaw movements.

[0115] In response to detecting the change 642 to the head pose information associated with the user 150 in FIG. 6D, the electronic device 120 presents, on the display 122, a 3D environment 602E associated with the first portion of the guided stretching session at time T5 in FIG. 6E. In response to detecting the change 642 to the head pose information associated with the user 150 in FIG. 6D and in accordance with a determination that the change 642 to the head pose information associated with the user 150 satisfies a criterion associated with a final visual cue 614C among the first plurality of visual cues, the electronic device 120 presents a notification 650 within the 3D environment 602E (e.g., visual feedback) indicating that the visual cue 614D has been completed for the first portion of the guided stretching session and that the first portion of the guided stretching session has been completed. In some implementations, the visual feedback is replaced with or supplemented by haptic feedback, audio feedback, spatial audio feedback, and/or the like as described above with reference to FIG. 6B.

[0116] As shown in FIG. 6E, the 3D environment 602E does not include any visual cues and includes the directional indicator 612D at the former location of visual cue 614C. As shown in FIG. 6E, the directional indicator 612D includes a visual representation (e.g., the number “0”) associated with a count of the first plurality of visual cues that have not been completed for the first portion of the guided stretching session at time T5.

[0117] In response to detecting the change 642 to the head pose information associated with the user 150 in FIG. 6D and in accordance with a determination that the change 642 to the head pose information associated with the user 150 satisfies the criterion associated with the final visual cue 614C among the first plurality of visual cues, the electronic device 120 presents, on the display 122, a 3D environment 672 associated with a second portion of the guided stretching session at time T6 in FIG. 6F. Continuing with this example, while presenting the 3D environment 672, the electronic device 120 presents a notification 660 instructing the user to follow the visual cues for the second portion of the guided stretching session by rotating their head from right-to-left by 180°.

[0118] As shown in FIG. 6F, the 3D environment 672 includes a second plurality of visual cues (e.g., detents) 664A, 664B, 664C, 664D, 664E, 664F for the second portion of the guided stretching session (e.g., a 180° rotational yaw movement of the user’s head). For example, the second plurality of visual cues 664A, 664B, 664C, 664D, 664E, 664F is evenly spaced along a 180° arc relative to the user’s current location. For example, if the second plurality of visual cues includes six visual cues, the electronic device 120 presents the six visual cues on the 180° arc with 30° between each of the six visual cues. In some implementations, the second plurality of visual cues 664A, 664B, 664C,

664D, 664E, 664F have similar sizes but appear different based on their depth and/or distance relative to the user’s current location.

[0119] According to some implementations, the electronic device 120 determines locations for the second plurality of visual cues 664A, 664B, 664C, 664D, 664E, 664F within the 3D environment 672 based on the user profile information and the head pose information associated with the user. In some implementations, the electronic device 120 determines the locations for the second plurality of visual cues 664A, 664B, 664C, 664D, 664E, 664F to suit the user’s height and/or sightline. In some implementations, the electronic device 120 determines the locations for the second plurality of visual cues 664A, 664B, 664C, 664D, 664E, 664F so as not to occlude or collide with objects within the 3D environment. In some implementations, the electronic device 120 determines the locations for the second plurality of visual cues 664A, 664B, 664C, 664D, 664E, 664F based on user flexibility/range of motion (taking user history into account) or user preferences.

[0120] As shown in FIG. 6F, the 3D environment 672 also includes a directional indicator 662 associated with a second stretch direction for the second portion of a guided stretching session (e.g., right-to-left or counterclockwise rotational movement). In one example, the electronic device 120 presents the directional indicator 662 at a deterministic origin location associated with the former location of the visual cue 614C when starting the second portion of the guided stretching session. In one example, the electronic device 120 presents the directional indicator 662 at a deterministic origin (default) location when starting the second portion of the guided stretching session (e.g., the location associated with the directional indicator 612A in FIG. 6B). In yet another example, the electronic device 120 presents the directional indicator 662 at a non-deterministic origin location based on the user’s current height and/or sightline. As shown in FIG. 6F, the directional indicator 662 includes a visual representation (e.g., the number “6”) associated with a count of the second plurality of visual cues that have not been completed for the second portion of the guided stretching session at time T6.

[0121] FIGS. 7A-7D illustrate a flowchart representation of a method 700 of presenting a guided stretching session with some implementations. In various implementations, the method 700 is performed at a computing system including non-transitory memory and one or more processors, wherein the computing system is communicatively coupled to a display device and one or more input devices (e.g., the electronic device 120 shown in FIGS. 1 and 3; the controller 110 in FIGS. 1 and 2; or a suitable combination thereof). In some implementations, the method 700 is performed by processing logic, including hardware, firmware, software, or a combination thereof. In some implementations, the method 700 is performed by a processor executing code stored in a non-transitory computer-readable medium (e.g., a memory). In some implementations, the computing system corresponds to one of a tablet, a laptop, a mobile phone, a near-eye system, a wearable computing device, or the like. In some implementations, the one or more input devices correspond to a computer vision (CV) engine that uses an image stream from one or more exterior-facing image sensors, a finger/hand/extremity tracking engine, an eye tracking engine, a touch-sensitive surface, one or more microphones, and/or the like.

[0122] As discussed above, many persons (e.g., the user 150 in FIG. 1) may spend a significant number of hours at their computers or other devices during both work and non-work hours. This time spent using a computer or other devices may negatively impact the posture of said person. As such, described herein is a method and device for promoting posture awareness and reducing an accumulated strain value for the user by presenting a guided stretching session.

[0123] As represented by block 702, while presenting a three-dimensional (3D) environment via the display device, the method 700 includes obtaining (e.g., receiving, retrieving, determining/generating, etc.) user profile information and head pose information for a user associated with the computing system. For example, the user profile information includes visual acuity metrics for the user, user height, user preferences, and/or the like. In some implementations, the head pose information includes 3DOF rotational values. In some implementations, the computing system also obtains FOV information indicating a current viewing frustum of the user based on the gaze direction.

[0124] In some implementations, the computing system or a component thereof (e.g., the interactive stretching engine 468A in FIG. 4E) obtains the user profile information 473 associated with the user. In some implementations, the computing system or a component thereof (e.g., the interactive stretching engine 468A in FIG. 4E) obtains the characterization vector 419 including head pose information and/or body pose information associated with the user.

[0125] As one example, with reference to FIG. 6A, the electronic device 120 presents, on the display 122, a 3D environment 602A associated with the guided stretching session at time T1. Continuing with this example, while presenting the 3D environment 602A, the electronic device 120 presents a notification 610 welcoming the user to the guided stretching session. In some implementations, the computing system or a component thereof (e.g., the rendering engine 550 in FIG. 5) renders the 3D environment (e.g., an XR environment) by compositing rendered XR content with a processed image stream of the physical environment 105 to produce rendered image frames of the 3D environment. In various implementations, the computing system or a component thereof (e.g., the presenter 560 in FIG. 5) presents the rendered image frames of the 3D environment to the user 150 via the one or more displays 312. One of ordinary skill in the art will appreciate that the aforementioned compositing process may not be applicable for fully virtual environments or optical see-through scenarios.

[0126] In some implementations, the display device corresponds to a transparent lens assembly, and wherein presenting the 3D environment includes projecting a least a portion of the 3D environment onto the transparent lens assembly. In some implementations, the display device corresponds to a near-eye system, and wherein presenting the 3D environment includes compositing at least a portion of the 3D environment with one or more images of a physical environment captured by an exterior-facing image sensor.

[0127] As represented by block 704, the method 700 includes determining locations for a first plurality of visual cues (e.g., detents) within the 3D environment for a first portion of a guided stretching session based on the user profile information and the head pose information, wherein the first portion of the guided stretching session corresponds to a first stretch direction. In some implementations, the

computing system or a component thereof (e.g., the interactive stretching engine 468A in FIG. 4E) determines the locations for the first plurality of visual cues (e.g., the stretching session visualizations 469 in FIG. 4E) based at least in part on the user profile information 473 associated with the user and the characterization vector 419 including head pose information and/or body pose information associated with the user

[0128] In some implementations, the computing system or a component thereof (e.g., the interactive stretching engine 468A in FIG. 4E) generates the first plurality of visual cues (e.g., the stretching session visualizations 469 in FIG. 4E) and determines locations for the first plurality of visual cues in response to detecting a user input to initiate a guided stretching session (e.g., a touch input, a hand/extremity tracking input, an eye tracking input, a voice command, or the like). In some implementations, the computing system or a component thereof (e.g., the interactive stretching engine 468A in FIG. 4E) generates the first plurality of visual cues (e.g., the stretching session visualizations 469 in FIG. 4E) and determines the locations for the first plurality of visual cues in response to detecting an accumulated strain value greater than a threshold value. In some implementations, the threshold value corresponds to a deterministic or non-deterministic value. According to some implementations, the threshold value corresponds to the fourth posture awareness threshold mentioned in provisional patent application No xxx,xxx, filed on xxx (Attorney Docket Number 27753-50497US1), which is incorporated by reference in its entirety.

[0129] In some implementations, as represented by block 706, the user profile information includes at least one of range of motion information or flexibility information associated with the user, and wherein spacing between the determined locations for the first plurality of visual cues is based on the range of motion information or the flexibility information associated with the user.

[0130] In some implementations, the computing system determines the locations for the first plurality of visual cues to suit the user's height and/or sightline. In some implementations, the computing system determines the locations for the first plurality of visual cues so as not to occlude or collide with objects within the 3D environment. In some implementations, the computing system determines the locations for the first plurality of visual cues based on user flexibility/range of motion (taking user history into account) or user preferences. In some implementations, the computing system tailors the guided stretching session based on the user's current and/or historical flexibility/range of motion. As one example, if the user is unable to complete the first portion of the guided stretching session, the computing system may decrease the range of motion for the second portion of the guided stretching session or for subsequent guided stretching sessions.

[0131] In some implementations, as represented by block 708, the method 700 includes obtaining (e.g., receiving, retrieving, or determining/generating) environment information associated with a physical environment including dimensional information for the physical environment, localization information for physical objects within the physical environment, and dimensional information for the physical objects within the physical environment, wherein the locations for the first plurality of visual cues are determined based on the user profile information, the head pose

information, and the environment information. In some implementations, the computing system or a component thereof (e.g., the interactive stretching engine **468A** in FIG. **4E**) obtains the context information vector **470** including environment state information **472** such as semantic labels for physical objects within the physical environment **105**, locations for physical objects within the physical environment **105**, and/or the like. In some implementations, the computing system or a component thereof (e.g., the interactive stretching engine **468A** in FIG. **4E**) determines the locations for the first plurality of visual cues based at least in part on the environment state information **472** so as not to occlude or collide with objects within the 3D environment.

[**0132**] For example, the environment information corresponds to SLAM information for localizing the user relative to the physical environment and determining the dimensions thereof. For example, the environment information corresponds to semantic segmentation information (e.g., labels) and/or the like. in order to identify physical objects within the physical environment as well as their dimensions and locations within the physical environment.

[**0133**] As represented by block **710**, the method **700** includes presenting, via the display device, the first plurality of visual cues for the first portion of the guided stretching session at the determined locations within the 3D environment and a directional indicator (e.g., an arrow, text indicating a direction, or the like). For example, with reference to FIG. **6A**, the electronic device **120** presents the 3D environment **602A** including a first plurality of visual cues (e.g., detents) **614A**, **614B**, **614C** for a first portion of a guided stretching session (e.g., a 90° rotational yaw movement of the user's head). As shown in FIG. **6A**, the 3D environment **602A** also includes a directional indicator **612A** associated with a first stretch direction for the first portion of a guided stretching session (e.g., left-to-right or clockwise rotational movement).

[**0134**] In some implementations, the first plurality of visual cues corresponds to volumetric extended reality (XR) objects. In some implementations, the first plurality of visual cues is overlaid on a representation of a physical environment while presented within the 3D environment. In some implementations, the first plurality of visual cues is composited with a representation of a physical environment while presented within the 3D environment.

[**0135**] In some implementations, as represented by block **712**, the directional indicator includes a visual representation (e.g., a numeral or an alphanumeric text string) associated with a count of the first plurality of visual cues that have not been completed for the first portion of the guided stretching session. As one example, with reference to FIG. **6A**, the directional indicator **612A** includes a visual representation (e.g., the number "3") associated with a count of the first plurality of visual cues that have not been completed for the first portion of the guided stretching session at time **T1**.

[**0136**] In some implementations, as represented by block **714**, the directional indicator is initially presented at an origin location relative to the head pose information associated with the user, and wherein the first plurality of visual cues is spaced in a 90-degree arc in the first stretch direction relative to the origin location. In some implementations, the first plurality of visual cues is evenly spaced along the 90°

arc. For example, if the first plurality of visual cues includes **N** visual cues, the computing system presents the **N** visual cues on the 90° arc with

$$\frac{90}{N}$$

degrees between each of the **N** visual cues.

[**0137**] In some implementations, the computing system initially presents the directional indicator at a deterministic or non-deterministic origin location. In one example, with reference to FIG. **6A**, the electronic device **120** presents the directional indicator **612A** at a deterministic (default or predefined) origin location when starting the guided stretching session. In another example, with reference to FIG. **6A**, the electronic device **120** presents the directional indicator **612A** at a non-deterministic origin location based on the user's current height and/or sightline. For example, the origin location is centered within the user's initial FOV.

[**0138**] As represented by block **716**, the method **700** includes detecting, via the one or more input devices, a change to the head pose information associated with the user. As one example, with reference to FIG. **6B**, while presenting the 3D environment **602B**, the electronic device **120** detects a change **622** to the head pose information associated with the user **150** such as a 30° clockwise rotational yaw movement.

[**0139**] As represented by block **718**, in response to detecting the change to the head pose information associated with the user, the method **700** includes: updating a location for the directional indicator based on the change to the head pose information associated with the user; and in accordance with a determination that the change to the head pose information associated with the user satisfies a criterion associated with a first visual cue among the first plurality of visual cues, providing at least one of audio, spatial audio, haptic, or visual feedback indicating that the first visual cue among the first plurality of visual cues has been completed for the first portion of the guided stretching session. In some implementations, the computing system or a component thereof (e.g., the interactive stretching engine **468A** in FIG. **4E**) updates the location of the directional indicator to track changes to the head pose information and/or the body pose information. In some implementations, the computing system or a component thereof (e.g., the interactive stretching engine **468A** in FIG. **4E**) also generates and provides at least one of audio, haptic, or visual feedback in accordance with a determination that the change to the head pose information and/or the body pose information associated with the user satisfies a criterion associated with a first visual cue among the first plurality of visual cues. In some implementations, the computing system or a component thereof (e.g., the strain decrease logic **465B** in FIG. **4C**) decreases the accumulated strain values and/or one or more muscle strain values in accordance with a determination that the change to the head pose information and/or the body pose information associated with the user satisfies a criterion associated with a first visual cue among the first plurality of visual cues.

[**0140**] For example, in response to detecting the change **622** to the head pose information associated with the user **150** in FIG. **6B**, the electronic device **120** presents the directional indicator **612B** at the former location of the visual cue **614A** in FIG. **6C**. As shown in FIG. **6C**, the



directional indicator **612B** includes a visual representation (e.g., the number “2”) associated with a count of the first plurality of visual cues that have not been completed for the first portion of the guided stretching session at time **T3**. Continuing with this example, in response to detecting the change **622** to the head pose information associated with the user **150** in FIG. **6B** and in accordance with a determination that the change **622** to the head pose information associated with the user **150** satisfies a criterion associated with the visual cue **614A**, the electronic device **120** presents a notification **630** within the 3D environment **602C** (e.g., visual feedback) in FIG. **6C** indicating that the visual cue **614A** has been completed for the first portion of the guided stretching session and to continue with the same motion.

[0141] In some implementations, the criterion associated with a visual cue is satisfied when the location for the directional indicator coincides with the location for the visual cue. In some implementations, the visual feedback is replaced with or supplemented by haptic feedback, audio feedback, and/or the like. According to some implementations, each of the visual cues, save for the final visual cue, is associated with similar haptic and/or audio feedback, and the final visual cue is associated with distinct haptic and/or audio feedback with a greater intensity, length, pitch, frequency, volume, and/or the like.

[0142] In some implementations, as represented by block **720**, the method **700** includes obtaining (e.g., receiving, retrieving, determining/generating, etc.) body pose information for the user associated with the computing system, wherein updating the location for the directional indicator is based on the change to at least one of the head pose information or the body pose information associated with the user. In some implementations, the computing system or a component thereof (e.g., the interactive stretching engine **468A** in FIG. **4E**) obtains the characterization vector **419** including head pose information and/or body pose information associated with the user. According to some implementations, the directional indicator tracks the head pose and/or body pose associated with the user.

[0143] In some implementations, as represented by block **722**, the method **700** includes: while presenting the first plurality of visual cues for the first portion of a guided stretching session at the determined locations within the 3D environment and the directional indicator, concurrently presenting, via the display device, a representation of the user within the 3D environment; and in response to detecting the change to the head pose information associated with the user, updating the representation of the user based on the change to the head pose information associated with the user. In some implementations, the computing system or a component thereof (e.g., the interactive stretching engine **468A** in FIG. **4E**) presents a representation of the user **150** within the 3D environment (e.g., associated with the stretching session visualizations **469** in FIG. **4E**) that tracks the changes to the head pose information and/or the body pose information associated with the user **150**. For example, the representation of the user **150** corresponds to an avatar with a face and/or body modeled on the user **150**.

[0144] In some implementations, as represented by block **724**, the feedback indicating that the first visual cue among the first plurality of visual cues has been completed for the first portion of the guided stretching session corresponds to spatial audio feedback spatialized based on the relative spacing between the representation of the user and the first

visual cue. As one example, when the first visual cue directly above the avatar’s head is completed, the computing system provides (e.g., outputs or plays) the spatial audio directly above the user’s head instead of emanating from the location of the first visual cue in front of the user.

[0145] According to some implementations, audio feedback may be spatialized to coincide with the location of the visual cue or at a distance in a direction of the stretch that satisfies the visual cue. In one example, with reference to yaw rotational movement, the audio feedback may be spatialized at a distance and at the angle such that the user moves their head side-to-side about they axis (e.g., based on the user’s head and sightline) to satisfy the current portion of the guided stretching session. In another example, with reference to roll rotational movement for the guided stretching session, the audio feedback may be spatialized relative to a location above the user’s head and at an angle such that the user tilts their head about the z axis (e.g., based on the user’s head and sightline) to satisfy the current portion of the guided stretching session. In yet another example, with reference to pitch rotational movement, the audio feedback may be spatialized at a location in front of the user and at an angle such that the user moves their head about the x axis (e.g., based on the user’s head and sightline) to satisfy the current portion of the guided stretching session.

[0146] In some implementations, as represented by block **726**, in accordance with the determination that the change to the head pose information associated with the user satisfies the criterion associated with the first visual cue among the first plurality of visual cues, the method **700** includes ceasing display of the first visual cue within the 3D environment. In response to detecting the change **622** to the head pose information associated with the user **150** in FIG. **6B**, the electronic device **120** presents the 3D environment **602C** in FIG. **6C**, including the visual cues **614B** and **614C** and the directional indicator **612B** at the former location of the visual cue **614A** without presenting the visual cue **614A**.

[0147] In some implementations, as represented by block **728**, in accordance with the determination that the change to the head pose information associated with the user satisfies the criterion associated with the first visual cue among the first plurality of visual cues, the method **700** includes maintaining display of the first visual cue within the 3D environment.

[0148] In some implementations, as represented by block **730**, in accordance with the determination that the change to the head pose information associated with the user satisfies the criterion associated with the first visual cue among the first plurality of visual cues, the method **700** includes changing an appearance of the first visual cue within the 3D environment. In some implementations, the computing system changes the color, texture, shape, etc. of the first visual target changes once reached. In some implementations, the computing system changes the opacity or translucency of the first visual target changes once reached. In some implementations, the computing system changes the first visual target to a dotted or ghosted outline representation once reached.

[0149] In some implementations, as represented by block **732**, in accordance with a determination that the change to the head pose information associated with the user satisfies a criterion associated with a final visual cue among the first plurality of visual cues, the method **700** includes providing at least one of audio, haptic, or visual feedback indicating that the final visual cue among the first plurality of visual

cues has been completed for the first portion of the guided stretching session; and ceasing display of the first plurality of visual cues. As represented by block 732, the method 700 further includes: after ceasing display of the first plurality of visual cues, determining locations for a second plurality of visual cues (e.g., detents) within the 3D environment for a second portion of a guided stretching session based on the user profile information and the head pose information, wherein the second portion of the guided stretching session corresponds to a second stretch direction; and concurrently presenting, via the display device, the second plurality of visual cues for the second portion of a guided stretching session at the determined locations within the 3D environment and the directional indicator. In some implementations, the computing system or a component thereof (e.g., the strain decrease logic 465B in FIG. 4C) decreases the accumulated strain values and/or one or more muscle strain values in accordance with a determination that the first portion of the guided stretching session has been completed. In some implementations, the computing system prompts the user to perform multiple repetitions of the first and/or second portions of the guided stretching session.

[0150] As one example, in response to detecting the change 642 to the head pose information associated with the user 150 in FIG. 6D and in accordance with a determination that the change 642 to the head pose information associated with the user 150 satisfies the criterion associated with the final visual cue 614C among the first plurality of visual cues, the electronic device 120 presents, on the display 122, a 3D environment 672 in FIG. 6F associated with a second portion of the guided stretching session at time T6. Continuing with this example, the 3D environment 672 in FIG. 6F includes a second plurality of visual cues (e.g., detents) 664A, 664B, 664C, 664D, 664E, 664F for the second portion of the guided stretching session (e.g., a 180° rotational yaw movement of the user's head). Further continuing with this example, the 3D environment 672 in FIG. 6F also includes a directional indicator 662 associated with a second stretch direction for the second portion of a guided stretching session (e.g., right-to-left or counterclockwise rotational movement).

[0151] For example, the second stretch direction corresponds to rotational movement opposite the first stretch direction. In some implementations, while presenting the second plurality of visual cues for the second portion of a guided stretching session, the directional indicator is presented at the location associated with the final visual cue among the first plurality of visual cues, and the second plurality of visual cues induces 180° of rotational movement in the second stretch direction relative to the stopping point for the first portion of the guided stretching session (e.g., the second plurality of visual cues is evenly spaced along a 180° arc in the second stretch direction).

[0152] In some implementations, while presenting the second plurality of visual cues for the second portion of a guided stretching session, the directional indicator is presented back at the origin location, and the second plurality of visual cues induces 90° of rotational movement in the second direction relative to the origin location (e.g., the second plurality of visual cues is evenly spaced along a 90° arc in the second stretch direction).

[0153] In some implementations, after completing the second portion of a guided stretching session, the computing system may induce the user into performing multiple rep-

etitions of the first and second portions of the guided stretching session. In some implementations, after completing the second portion of a guided stretching session, the computing system may induce the user into performing a third portion of the guided stretching session in a third stretch direction orthogonal to the first and second stretch direction. As one example, the first stretch direction corresponds to left-to-right rotational yaw movement (about the y-axis of the user's head), the second stretch direction corresponds to right-to-left rotational yaw movement (about the y-axis of the user's head), the third stretch direction corresponds to left-to-right roll rotational movement (about the z-axis of the user's head), or the like. One of ordinary skill will appreciate that the guided stretching session may induce the user to perform any of the roll, pitch, and/or rotational yaw movements with their head and/or body in an arbitrary sequence in various implementations. In some implementations, the computing system prompts the user to perform different stretch types during the guided stretching session such as a shoulder shrug or the like in addition to head rotation.

[0154] In some implementations, as represented by block 734, in accordance with a determination that the change to the head pose information associated with the user does not satisfy the criterion associated with the first visual cue among the first plurality of visual cues (within a deterministic or non-deterministic length of time), the method 700 includes providing one or more hints to the user associated with the first portion of the guided stretching session. For example, the one or more hints correspond to spatial audio feedback, haptic feedback, text feedback, illustrations, arrows and/or other symbols, audio commands, and/or the like to aid the user in reaching the first visual cue among the first plurality of visual cues.

[0155] While various aspects of implementations within the scope of the appended claims are described above, it should be apparent that the various features of implementations described above may be embodied in a wide variety of forms and that any specific structure and/or function described above is merely illustrative. Based on the present disclosure one skilled in the art should appreciate that an aspect described herein may be implemented independently of any other aspects and that two or more of these aspects may be combined in various ways. For example, an apparatus may be implemented and/or a method may be practiced using any number of the aspects set forth herein. In addition, such an apparatus may be implemented and/or such a method may be practiced using other structure and/or functionality in addition to or other than one or more of the aspects set forth herein.

[0156] It will also be understood that, although the terms "first", "second", etc. may be used herein to describe various elements, these elements should not be limited by these terms. These terms are only used to distinguish one element from another. For example, a first media item could be termed a second media item, and, similarly, a second media item could be termed a first media item, which changing the meaning of the description, so long as the occurrences of the "first media item" are renamed consistently and the occurrences of the "second media item" are renamed consistently. The first media item and the second media item are both media items, but they are not the same media item.

[0157] The terminology used herein is for the purpose of describing particular implementations only and is not

intended to be limiting of the claims. As used in the description of the implementations and the appended claims, the singular forms “a”, “an”, and “the” are intended to include the plural forms as well, unless the context clearly indicates otherwise. It will also be understood that the term “and/or” as used herein refers to and encompasses any and all possible combinations of one or more of the associated listed items. It will be further understood that the terms “comprises” and/or “comprising,” when used in this specification, specify the presence of stated features, integers, steps, operations, elements, and/or components, but do not preclude the presence or addition of one or more other features, integers, steps, operations, elements, components, and/or groups thereof.

**[0158]** As used herein, the term “if” may be construed to mean “when” or “upon” or “in response to determining” or “in accordance with a determination” or “in response to detecting,” that a stated condition precedent is true, depending on the context. Similarly, the phrase “if it is determined [that a stated condition precedent is true]” or “if [a stated condition precedent is true]” or “when [a stated condition precedent is true]” may be construed to mean “upon determining” or “in response to determining” or “in accordance with a determination” or “upon detecting” or “in response to detecting” that the stated condition precedent is true, depending on the context.

What is claimed is:

**1.** A method comprising:

at a computing system including non-transitory memory and one or more processors, wherein the computing system is communicatively coupled to a display device and one or more input devices via a communication interface:

while presenting a three-dimensional (3D) environment via the display device, obtaining user profile information and head pose information for a user associated with the computing system;

determining locations for a first plurality of visual cues within the 3D environment for a first portion of a guided stretching session based on the user profile information and the head pose information, wherein the first portion of the guided stretching session corresponds to a first stretch direction;

presenting, via the display device, the first plurality of visual cues for the first portion of the guided stretching session at the determined locations within the 3D environment and a directional indicator;

detecting, via the one or more input devices, a change to the head pose information associated with the user; and

in response to detecting the change to the head pose information associated with the user:

updating a location for the directional indicator based on the change to the head pose information associated with the user; and

in accordance with a determination that the change to the head pose information associated with the user satisfies a criterion associated with a first visual cue among the first plurality of visual cues, providing at least one of audio, haptic, or visual feedback indicating that the first visual cue among the first plurality of visual cues has been completed for the first portion of the guided stretching session.

**2.** The method of claim **1**, further comprising:

obtaining environment information associated with a physical environment including dimensional information for the physical environment, localization information for physical objects within the physical environment, and dimensional information for the physical objects within the physical environment, wherein the locations for the first plurality of visual cues are determined based on the user profile information, the head pose information, and the environment information.

**3.** The method of claim **1**, further comprising:

obtaining body pose information for the user associated with the computing system, wherein updating the location for the directional indicator is based on the change to at least one of the head pose information or the body pose information associated with the user.

**4.** The method of claim **1**, wherein the directional indicator includes a visual representation associated with a count of the first plurality of visual cues that have not been completed for the first portion of the guided stretching session.

**5.** The method of claim **1**, further comprising:

while presenting the first plurality of visual cues for the first portion of a guided stretching session at the determined locations within the 3D environment and the directional indicator, concurrently presenting, via the display device, a representation of the user within the 3D environment; and

in response to detecting the change to the head pose information associated with the user, updating the representation of the user based on the change to the head pose information associated with the user.

**6.** The method of claim **5**, wherein the feedback indicating that the first visual cue among the first plurality of visual cues has been completed for the first portion of the guided stretching session corresponds to spatial audio feedback spatialized based on the relative spacing between the representation of the user and the first visual cue.

**7.** The method of claim **1**, wherein the user profile information includes at least one of range of motion information or flexibility information associated with the user, and wherein spacing between the determined locations for the first plurality of visual cues is based on the range of motion information or the flexibility information associated with the user.

**8.** The method of claim **1**, wherein the directional indicator is initially presented at an origin location relative to the head pose information associated with the user, and wherein the first plurality of visual cues is spaced in a 90-degree arc in the first stretch direction relative to the origin location.

**9.** The method of claim **1**, further comprising:

in accordance with a determination that the change to the head pose information associated with the user does not satisfy the criterion associated with the first visual cue among the first plurality of visual cues, providing one or more hints to the user associated with the first portion of the guided stretching session.

**10.** The method of claim **1**, further comprising:

in accordance with the determination that the change to the head pose information associated with the user satisfies the criterion associated with the first visual cue among the first plurality of visual cues, ceasing display of the first visual cue within the 3D environment.

- 11.** The method of claim 1, further comprising:  
in accordance with the determination that the change to the head pose information associated with the user satisfies the criterion associated with the first visual cue among the first plurality of visual cues, maintaining display of the first visual cue within the 3D environment.
- 12.** The method of claim 1, further comprising:  
in accordance with the determination that the change to the head pose information associated with the user satisfies the criterion associated with the first visual cue among the first plurality of visual cues, changing an appearance of the first visual cue within the 3D environment.
- 13.** The method of claim 1, wherein the first plurality of visual cues corresponds to volumetric extended reality (XR) objects.
- 14.** The method of claim 1, wherein the first plurality of visual cues is overlaid on a representation of a physical environment while presented within the 3D environment.
- 15.** The method of claim 1, wherein the first plurality of visual cues is composited with a representation of a physical environment while presented within the 3D environment.
- 16.** The method of claim 1, further comprising:  
in accordance with a determination that the change to the head pose information associated with the user satisfies a criterion associated with a final visual cue among the first plurality of visual cues:  
providing at least one of audio, haptic, or visual feedback indicating that the final visual cue among the first plurality of visual cues has been completed for the first portion of the guided stretching session; and  
ceasing display of the first plurality of visual cues;  
after ceasing display of the first plurality of visual cues, determining locations for a second plurality of visual cues within the 3D environment for a second portion of a guided stretching session based on the user profile information and the head pose information, wherein the second portion of the guided stretching session corresponds to a second stretch direction; and  
concurrently presenting, via the display device, the second plurality of visual cues for the second portion of a guided stretching session at the determined locations within the 3D environment and the directional indicator.
- 17.** A device comprising:  
one or more processors;  
a non-transitory memory;  
an interface for communicating with a display device and one or more input devices; and  
one or more programs stored in the non-transitory memory, which, when executed by the one or more processors, cause the device to:  
while presenting a three-dimensional (3D) environment via the display device, obtain user profile information and head pose information for a user associated with the computing system;  
determine locations for a first plurality of visual cues within the 3D environment for a first portion of a guided stretching session based on the user profile information and the head pose information, wherein the first portion of the guided stretching session corresponds to a first stretch direction;  
present, via the display device, the first plurality of visual cues for the first portion of the guided stretching session at the determined locations within the 3D environment and a directional indicator;  
detect, via the one or more input devices, a change to the head pose information associated with the user; and  
in response to detecting the change to the head pose information associated with the user:  
present, via the display device, the first plurality of visual cues for the first portion of the guided stretching session at the determined locations within the 3D environment and a directional indicator;  
detect, via the one or more input devices, a change to the head pose information associated with the user; and  
in response to detecting the change to the head pose information associated with the user:  
update a location for the directional indicator based on the change to the head pose information associated with the user; and  
in accordance with a determination that the change to the head pose information associated with the user satisfies a criterion associated with a first visual cue among the first plurality of visual cues, provide at least one of audio, haptic, or visual feedback indicating that the first visual cue among the first plurality of visual cues has been completed for the first portion of the guided stretching session.
- 18.** The device of claim 17, wherein the directional indicator includes a visual representation associated with a count of the first plurality of visual cues that have not been completed for the first portion of the guided stretching session.
- 19.** The device of claim 17, wherein the user profile information includes at least one of range of motion information or flexibility information associated with the user, and wherein spacing between the determined locations for the first plurality of visual cues is based on the range of motion information or the flexibility information associated with the user.
- 20.** The device of claim 17, wherein the one or more programs further cause the device to:  
in accordance with a determination that the change to the head pose information associated with the user does not satisfy the criterion associated with the first visual cue among the first plurality of visual cues, provide one or more hints to the user associated with the first portion of the guided stretching session.
- 21.** A non-transitory memory storing one or more programs, which, when executed by one or more processors of a device with an interface for communicating with a display device and one or more input devices, cause the device to:  
while presenting a three-dimensional (3D) environment via the display device, obtain user profile information and head pose information for a user associated with the computing system;  
determine locations for a first plurality of visual cues within the 3D environment for a first portion of a guided stretching session based on the user profile information and the head pose information, wherein the first portion of the guided stretching session corresponds to a first stretch direction;  
present, via the display device, the first plurality of visual cues for the first portion of the guided stretching session at the determined locations within the 3D environment and a directional indicator;  
detect, via the one or more input devices, a change to the head pose information associated with the user; and  
in response to detecting the change to the head pose information associated with the user:

update a location for the directional indicator based on the change to the head pose information associated with the user; and

in accordance with a determination that the change to the head pose information associated with the user satisfies a criterion associated with a first visual cue among the first plurality of visual cues, provide at least one of audio, haptic, or visual feedback indicating that the first visual cue among the first plurality of visual cues has been completed for the first portion of the guided stretching session.

**22.** The non-transitory memory of claim **21**, wherein the directional indicator includes a visual representation associated with a count of the first plurality of visual cues that have not been completed for the first portion of the guided stretching session.

**23.** The non-transitory memory of claim **21**, wherein the user profile information includes at least one of range of motion information or flexibility information associated with the user, and wherein spacing between the determined locations for the first plurality of visual cues is based on the range of motion information or the flexibility information associated with the user.

**24.** The non-transitory memory of claim **21**, wherein the one or more programs further cause the device to:

in accordance with a determination that the change to the head pose information associated with the user does not satisfy the criterion associated with the first visual cue among the first plurality of visual cues, provide one or more hints to the user associated with the first portion of the guided stretching session.

\* \* \* \* \*